SECRECY VERSUS OPENNESS

INTERNET SECURITY AND THE LIMITS OF OPEN SOURCE AND PEER PRODUCTION

ANDREAS SCHMIDT

Open source and peer production have been praised as organisational models that could change the world for the better. It is commonly asserted that almost any societal activity could benefit from distributed, bottom-up collaboration — by making societal interaction more open, more social, and more democratic. However, we also need to be mindful of the limits of these models. How could they function in environments hostile to openness? Security is a societal domain more prone to secrecy than any other, except perhaps for romantic love. In light of the destructive capacity of contemporary cyber attacks, how has the Internet survived without a comprehensive security infrastructure? *Secrecy versus Openness* describes the realities of Internet security production through the lenses of open source and peer production theories. The study offers a glimpse into the fascinating communities of technical experts, who played a pivotal role when the chips were down for the Internet after large-scale attacks. After an initial flirtation with openness in the early years, operational Internet security communities have put in place a form of social production that resembles the open source model in many aspects, but is substantially less open.

ISBN 978-90-8891-988-6

Image Copyright Rafal Olbinski, Courtesy Patinae, Inc. www.patinae.com

SECRECY VERSUS OPENNESS

INTERNET SECURITY AND THE LIMITS OF OPEN SOURCE AND PEER PRODUCTION

Proefschrift

ter verkrijging van de graad van doctor
aan de Technische Universiteit Delft,
op gezag van de Rector Magnificus prof. ir. K.C.A.M. Luyben,
voorzitter van het College voor Promoties,
in het openbaar te verdedigen op maandag 3 november 2014 om 10:00 uur
door

Andreas SCHMIDT

Magister Artium in Political Science,
Medieval and Modern History, Kassel University
geboren te Bad Arolsen, Duitsland.

Dit proefschrift is goedgekeurd door de promotor:
Prof.dr. M.L. Mueller

Samenstelling promotiecommissie:

Rector Magnificus,	voorzitter
Prof.dr. M.L. Mueller,	Technische Universiteit Delft, promotor
Prof.dr. K. Rannenberg,	Goethe-University Frankfurt
Prof.dr.ir. E. Huizer,	Universiteit Utrecht
Prof.dr. R. Deibert	Universiteit Toronto
Prof.dr. N.A.N.M. Eijk,	Vrije Universiteit Amsterdam
Prof.dr.ir. J. van den Berg,	Technische Universiteit Delft
Prof.dr. M.J.G. van Eeten,	Technische Universiteit Delft
Prof.dr.ir. M.F.W.H.A. Janssen,	Technische Universiteit Delft, reservelid

Distributed by
Delft University of Technology, Faculty of Technology, Policy and Management,
Jaffalaan 5, 2628BX Delft, The Netherlands, and
Andreas Schmidt, Laubestraße 26, 60594 Frankfurt am Main, Germany.

Printed in The Netherlands

Published by Uitgeverij BOXPress

ISBN 978-90-8891-988-6

Keywords: information security, peer production, open source, internet governance, security governance, incident response

The book has been composed in Adobe Caslon Pro and Baron Neue. The manuscript was written with Scrivener, aided by Sente, Devonthink, and Word.

Cover designed by Michaela Spohn, michaelaspohn.de.

The Marriage of Figaro by Rafal Olbinski as cover artwork.
© Rafal Olbinski, courtesy Patinae, Inc., www.patinae.com

Table of Contents

Acknowledgements

Good research requires various methods of efficient procrastination. A superb source of comfort for aspiring scholars certainly is phdcomics.com, a website offering wisdoms on the lives of *promovendi* that are as amusing as depressing. All too often, the brain is buzzing on other things than research, the fingers are resting on the keyboard without moving, no souffleuse would whisper words of academic wisdom and no siren lures one into the beauty of scholarly reason. Then, the adage "There is no team in thesis"[1] sounds too convincing. But only then. I'm deeply grateful to all the individuals without whom this project would have been hardly viable, more difficult, not as good and considerably less fun.

The mundane foundations of this research project were laid by two organisations. The *Next Generation Infrastructure Foundation* — and its financier, the Dutch taxpayers — have generously funded my research position at Delft University of Technology. The communications service provider XS4ALL facilitated my supervisor's professorship and stay here in the Netherlands. I am deeply grateful for this support.

The Faculty of Technology, Policy and Management and its section for Information and Communication Technology have served as my academic home base. It has been the warm, cordial, caring, and collegial place that allowed me to focus on my research. Harry Bouwman has helped with invaluable counsel during my research project. Yao-Hua Tan has been the kind of manager you need when balancing research with newly arrived private responsibilities requires a flexible response. Marijn Janssen has guided me safely through the final stages of my project. Jo-Ann and Eveline have spoiled me with office supplies, bureaucratic support, and

[1] J. Cham, "I in team", PhD comics, March 25, 2013, http://www.phdcomics.com/comics/archive.php?comicid=1570

most of all with their kindness and friendship. My colleagues Bram, Fatemeh, Jan, Janneke, Tineke, Jolien, Jos, Mark, Nitesh, Mohsen, Martijn, and Anne Fleur have all contributed to an enjoyable intellectual and social climate and supported my endeavour in various ways. Over the years, I have benefited enormously from inspiring discussions and pleasant conversations with Michel van Eeten and his team, including Shirin Tabatabaie, Floris Kreiken, and Hadi Asghari. Hadi and Coen van Leeuwen will have helped in the very last part of this endeavour.

This research project has benefited immensely from the insights shared by practicing experts of Internet security, who were willing to discuss, often for the better part of a day, any aspect of their subject area. I am deeply grateful to Jan Priisula, Gadi Evron, Toni Koivunen, Rain Ottis, Hillar Aarelaid, Mikko Hippönen, Gorazd Božič, Kaido Raiend, Kurtis Eric Lindquist, Bill Woodcock, Merike Kaeo, Kauto Huopio, Andrus Padar, David Watson, Niklas Schiffler, Richard Perletto, Steve Santorelli, Scott McIntyre, Paul Vixie, Andre DiMinio, John Crain, Chris Lee, Tillmann Werner, Rick Wesson, Jose Nazario, Heli Tiirmaa-Klaar, and all the other persons, who prefer to remain anonymous or who I talked to 'only' on an informal basis.

Furthermore, I feel obliged to a number of wonderful individuals, who invited me to intellectual exchanges, gave valuable input (some extensive, some brief), collaborated with me on papers and books, invited me to present my work, or commented theron: Brenden Kuerbis, Axel Arnbak, Ben Wagner Ingo Stützle, Roxana Radu, Jean-Marie Chenou, Markus Beckedahl, Kate Coyer, Stefaan Verhulst, Jason Healey, Jan-Frederik Kremer, Benedikt Müller, Min-Chun Ku, Christian Pentzold, Giampiero Giacomello, and many others. Ralf Bendrath deserves a special thank you for luring me into this project and our joint time in Delft.

The finalising of this endeavour was eased by a few helping hands. Deserving of thanks, Katarina has created the transcripts, Danielle and Patrick have proofread this study and provided helpful editorial tips. The stellar design of the cover is the result of Michaela's artistry. The cover has been made possible by the more than generous efforts of Sherri Nahan from Rafal Olbinski's agency Patinae Inc.

And then there is Milton Mueller. I could not have been happier with his support for my pursuits, his supervision of my project and his patience. Working with a man who knows his subject area like hardly any other has been a blessing and joy. He has helped me out endlessly with super-prompt, sharp and witty comments, admonishing music collections, and revisions of countless, increasingly less imperfect drafts. Imagine a chivalrous deep bow of warm-hearted gratitude for all of that, Milton. And for the fun we had during the project, a long-lasting smile.

My final words here are to express my gratitude and love for my friends and family. My friends have provided me with opportunities for amusement, retreats, and invaluable spaces for undisturbed writing whenever dearly needed. My siblings, my father, and my mother, who sadly passed :'(away this summer, and my in-laws have been more than supportive and encouraging. With the easy naturalness of a new arrival, Ada has turned things upside down with her impressive talent for creative sleeping and creating deeper-than-deep affection. Michaela, your ingenious cover is worth praise by itself. And yet, it has only been your smallest contribution to this project. Your sheer patience and warm-hearted support will always be remembered.

Andreas Schmidt

1 Introduction

"[W]e find ourselves in the midst of a battle
over the institutional ecology of the digital environment"
Yochai Benkler

"The digital is the realm of the open: open source, open resources, open doors.
Anything that attempts to close this space should be recognized for what it is: the enemy."
Digital Humanities Manifesto

"Loose lips sink ships"
World War II poster claim, USA[2]

The Internet has acquired the status of an indispensable infrastructure for a tremendous range of communication and transactions around the globe. The applications are interwoven in daily activities to such an extent that modern societies would be hamstrung without their ICT systems and the Internet up and running. With ever more social activities based on the Internet, the stakes are increasing for potential losses of Internet functionalities, as they would result in a substantial reduction of productivity, and harm our ability to communicate and share information efficiently. The vulnerabilities of networked computers are exposed in manifold Internet security problems such as spam, viruses, phishing, identity theft, botnets, civilian and state-sponsored denial of service attacks, and unauthorised intrusion into private networks. The scale of these problems is the subject of widespread discussion and debate. The very same security problems are at the core of five distinct, but interlinked discourses on cybercrime, cyberterror, cyberespionage, cyberwarfare, and critical infrastructure protection. While estimates about potential risks and real damages vary widely, it is safe to say that damages caused only by cybercrime amount to billions of US Dollars or Euros and thus provide strong incentives for mitigating activities (van Eeten & Bauer, 2008a; Anderson et al., 2012). Given the scale of the problem and its distributed, border-crossing nature,

[2] Benkler 2006, p. 2; UCLA Center for Digital Humanities 2008; "Loose Lips Sink Ships," 2001.

the question arises as to which governance and security production approaches are best suited to deal with these Internet security problems.

The Internet offers new opportunities for ‘creative destruction,’ thus endangering venerable entrenched interests and institutions. The transformational force of ICT and the Internet has substantially diminished the role of some traditional intermediaries, altered entire economic sectors, and given birth to entire new businesses, new economic players and thereby new political actors. Advertisement, distribution of informational goods such as software, music, and film; retail in general; libraries; storage and retrieval of written text, maps, images, or whatever kind of information; travel agencies; dating; journalism; public relations; interhuman communication; payment and banking — the list could well be extended by dozens of additional examples. The Internet has also left its mark in the political domain, altering political revolutionising, mass upheavals, intelligence, political campaigning, and the martial domain of “politics by other means” (Clausewitz). And yet, the impact and possibilities of the potentially transformative organisational changes on polity and politics are far from clear.

Enter peer production, the term that describes the particular features of a form of production that has been showcased in the creation of open source software like Linux or in collaborative content production projects like Wikipedia. Despite massive onslaughts by entrenched interests, defamatory attacks, and attempts to undermine the legal basis of open source projects, open source production still provides substantial products and services.

Recently though, the ideas of openness and free information exchange has come under pressure on various fronts. The hyperbolical idea of a ‘Twitter revolution’ has waned with the reactionary counters to the Arab Spring. Microsoft’s desktop monopoly is on the verge of being succeeded by the new personal computing duopoly of Apple’s even more closed and integrated products and Google’s pseudo-open Android. These developments might ring in a roll-back, rendering open production to an interim phenomenon, or they could be mere temporary setbacks.

The underlying questions which have driven this research project encompass the general limits of peer production, its applicability to domains other than software or encyclopaedia production. Security production, which usually comes with more or less secrecy, appears to be the most unlikely case to apply the pure form of peer production with all the openness it entails. Open collaboration contradicts secrecy, but the production of Internet security requires distributed collaboration. This thesis looks into the interplay of secrecy, openness, Internet security, and peer production.

1.1 Organising Internet security

The nature of Internet security is such that it requires governance and operational mechanisms that cross organisational and jurisdictional lines. Attacking machines, resources, and personnel are geographically distributed, exploiting vulnerabilities of resources belonging to different organisations in different countries with different laws.

These characteristics of Internet security problems raise the question about the best ways to organise the production of Internet security. The governance of security has not only a substantial effect on whether security is effectively and sufficiently taken care of. Given their often-clandestine modes of operating, traditional security institutions come with a price tag for democratic principles such as transparency and accountability. The governance of security can therefore have a significant impact on shared societal values. Modes of security governance can differ in a variety of ways, spanning from the degree of state involvement in policy formulation, policy implementation or security operations, the role of coercion, distribution of authority, internal hierarchies, to the role of private actors in the security architecture, the fora, techniques and depth of sharing information, the kind of threats to Internet security addressed to the kind of objects of Internet security dealt with by the governance form.

The political response to Internet-based threats, risks, and vulnerabilities has been a mixture of increasing public awareness, fostering private self-regulation or public-private cooperation. Other responses included creating Internet security groups within traditional state-based security organisations, supporting international incident-response exercises, setting up secretive monitoring programmes, and increasing military cyber-units. Consequently, the current organisational landscape of Internet security governance is characterised by a variety of governance forms. They range from international conventions (Council of Europe), initiatives launched by international organisations (ITU), regional directives (EU), unilateral hegemony (NSA-led monitoring system), to national and regional public-private partnerships — to name only a few. Emerging Internet security institutions have been organised along national lines, but also in terms of their subject matter. Different security problems like viruses, denial-of-service-attacks, botnets, spam, and phishing appear to be dealt with in different, yet occasionally overlapping organisational and institutional settings; one might call these "problem-specific governance islands". Finally, different sectors in the economy follow individual approaches to deal with Internet security problems.

The news headlines of major media outlets indicate that Internet security politics are more and more driven by the actors that have always played a crucial role in nation states' security politics: states, international organisations, police forces, military and intelligence agencies. In recent years, there has apparently been a trend to revert to hierarchical means of governance — be it national governments introducing harsher Internet legislation, surveillance, filtering or blocking measures, or interest groups trying to leverage ICANN's authoritative control over critical Internet resources. International organisations have raised their voices, asking to be granted more important roles in tackling Internet security problems. The political rhetoric accompanying discussions on Internet security often highlights an alleged lack of response capabilities, institutions, and organisations and consequential calls for the involvement of traditional security institutions (Dunn Cavelty 2013, 2007).

1.2 The rise of remote collaboration

One of the many changes the Internet has brought is a revolutionising of collaboration irrespective of the geographical location of the participants. Earlier generations of telecommunications like telegraphy, phone, fax, and telex facilitated distributed organisations with ways to coordinate their activities and exchange information. The Internet, however, has turned things upside down. New technologies in general allow for reorganising existing organisational, political and production processes (Fohler 2003; Singh 2002; Skolnikoff 1993) . And the Internet has substantially eased, if not enabled the formation and self-organisation of geographically widely dispersed teams; it allows for new ways of organising tasks and processes on any societal level. One of the most intriguing developments of recent years has been the emergence of distributed, self-organised, bottom-up governed production of informational artefacts.

Scientific endeavours to analyse and understand these changes have observed a variety of new or altered forms of organisation, labelled with a fleet of new concepts and models. "Distributed collaboration", "crowdsourcing", "distributed problem solving networks" and "collaborative communities" are but a few examples of the terms coined in this debate.[3]

[3] More on these concepts in chapter 2.

Some of the most elaborated reflections are provided by the body of literature on open source or peer production. The latter term can roughly be conceptualized as a distributed form of collaboration for creating goods by networks of actors who are relatively equal in their overall status. Peer and open source production have been identified as potentially one of the most disruptive innovations in modes of production (Benkler 2002, 2006; Weber 2000, 2004). Since the limits of this mode of production have yet to be explored, this potential could be both in practice and in theory. Weber characterised open source software projects as a new way of organising production. Absence of participation barriers, voluntariness as a founding principle, and sharing of production resources based on the open-source ideology and a certain property regime have founded a model of "distributed innovation" that can be more effective than other types of networked collaboration such as platforms or subcontracting (2004, pp. 231–243). Benkler additionally analysed projects beyond free and open source software (FOSS) and similarly identified "a new modality of organising production", which is built on reduced transaction costs facilitated by the rise of the Internet (Benkler 2006, p. 60). This form of production, termed "commons-based peer production", is characterised as "radically decentralised, collaborative, and nonproprietary; based on sharing resources and outputs among widely distributed, loosely connected individuals who cooperate with each other without relying on either market signals or managerial commands" (2006, p. 60).

For Weber and Benkler, the open source or peer form of production is generic and not bound to a specific good such as software or encyclopaedias. But are they applicable to the production of basically any informational good? What are the limits of peer production? The unclear boundaries of these new forms, which have been facilitated by the Internet, raise the question as to whether they can be applied to mitigate the very security concerns created on and by this new communicational infrastructure.

1.3 Existence and viability of peer security production

The idea of the peer production of Internet security extends two trends. The first one is the empowerment of individuals and self-organised groups by Internet-based collaboration. The second trend is rooted in studies of governance and security governance, i.e., the diversification of actors providing security away from a state-only system towards "hybridisation of governance" (Schuppert 2008, pp. 34-37).

Just as, say, the existence of private security organisations does not imply the absence of different organisational approaches like state-driven security organisations, peer production of Internet security would not imply that Internet security was entirely provided by such distributed, collaborative, self-organised efforts. The problem this study seeks to address is obviously not whether Internet security can be provided entirely by peer efforts. States and market actors will hardly be relegated to mere spectators in in the security arena. The governance of complex issues often happens in multiple fora, by large numbers of organisations and individuals. The existence of peer production in Internet security would not completely replace security provisioning by markets, hierarchies, money, and coercion.

The research problem rather is whether peer production of Internet security is viable at all. The existing literature is not specific about the limits of peer production. It likewise has few answers as to which phenotypes of organisation arise in an environment that on the one hand is as peer-production friendly as a network of technological Internet experts, but as hostile to openness as the secrecy-prone domain of security production.

1.3.1 Theoretical problem

The application of modes of peer production in Internet security, a domain that potentially requires secrecy, poses a theoretical and scientific problem largely unaddressed by existing theories and literatures. The organisation of security production usually includes areas of secrecy. Peer production and open source production are linked to the openness and accessibility of information and goods necessary for production. There is no developed model of organisational governance that would avoid the apparent contradiction between secrecy and openness. And empirical research on peer production in fields beyond software development and distributed web-based content creation is almost completely absent.

The theory and viability of peer and open source production has most thoroughly been explored by Benkler (2006) and Weber (2004). Weber's empirical analysis focuses on software production, Benkler's on Internet-based collaboration. Both have made first attempts to generalise peer and open source production as genuine modes of production applicable in any societal field; and both have taken initial steps towards an explanatory model and theoretical core. Yet, neither theoretical model clearly specifies the defining characteristics that constitute peer production nor do they sufficiently describe the circumstances under which it is likely to arise. While open source as a production model for software has been lucidly described by Weber, existing models of peer production do not offer sufficient explanations and predictions about the viability of peer production or provide guidance as to its conceivable mutations under varying conditions. In other words: existing theories do not cover all aspects of the phenomenon of peer production. They do not explain what might happen if one feature of peer production were missing, or whether, if certain features are missing, it can be classified as peer production at all.

The puzzle that has driven this research project is whether secrecy will make peer production completely inappropriate and require a return to familiar firm- or market-based modes of production. Alternatively, will it lead to new modes of production that resemble peer production, but cannot quite be classified as such, as they deviate in defining characteristics?

Existing bodies of literature have not had much to contribute to these problems. Traditional security studies, which are rooted in international relation studies and questions of war, military and alliances, have yet to digest the rise of non-state actors and hybrid public-private security arrangements. This branch of security deals with physical damages and very tangible means of security. The latter lack the characteristics of intangible informational goods — a feature that is deemed a prerequisite for peer production—: extremely low investment and low transaction costs.

Research on the economics of Internet security has gained momentum amongst Internet security researchers. Originally developed and applied mostly in information and computer science, the underlying idea of this field is that Internet security problems cannot be understood entirely in terms of technical insufficiencies and vulnerabilities. There are also economic incentives that cause actors in the field to behave in certain ways that create security problems. Hence, researchers using this approach search for solutions that take into account the economic incentives of actors as well as technical vulnerabilities. This approach has brought forward some fundamental and practically valuable insights (Anderson & Moore, 2006, 2007; Anderson, Böhme, Clayton, & Moore, 2009, 2008; Tsiakis & Sthephanides, 2005; Ioannidis, Pym, & Williams, 2009; Bauer, van Eeten, & Chattopadhyay, 2008; Bauer & van Eeten, 2009; van Eeten 2008; van Eeten & Bauer, 2009; Moore, Clayton, & Anderson, 2009). However, it has so far contributed little to organisational issues such as the problem of distributing and granting access to information in Internet security provisioning.[4]

[4] I use security provisioning and security production mostly synonymously in this thesis. As does Benkler, apparently: "'Provisioning' refers to efforts aimed at producing a particular good that would not otherwise exist. 'Allocating' refers to decisions about how a good that exists, but is scarce relative

Footnote continued on the next page.

Similarly, the body of literature on Internet governance studies lacks depth when it comes to role of technical communities and collaborative production of Internet security. In a discussion paper on the various governance issues of Internet governance, Bauer assumed that self-organisation in the area of "security of networks and information systems" would be "probably not feasible or not effective"; in the areas of "cyber security, cyber crime", since not enough previous information and research would exist (Bauer 2005, p. 17). In an attempt with a similar direction written at a similar time, Mathiason et al. catalogued the "state of play" of Internet governance. While they correctly stated that "informal non-state actors" provided Internet governance, no "informal non-state" actor was mentioned as a provider of security governance (Mathiason, Mueller, Klein, Holitscher, & McKnight, 2004). Existing literature has apparently little to contribute to the question of peer production of Internet security, neither empirically nor theoretically.

Next to — or possibly related to — the secrecy-vs.-openness problem is that distributed collaborative security production requires the sharing of potentially delicate information. Peer production would hence require a significant degree of trust among the contributors. The glue holding individuals together in "collaborative communities" (Adler & Heckscher, 2006), to which peer production is closely related, is trust. Trust had already been identified by Powell (1990) as the governing principle of networks, as opposed to hierarchies and markets. A common functionalist argument is that information exchange and collaboration is enabled by what is called "swift trust" (Adler & Heckscher, 2006; Osterloh & Rota, 2004), a mode of trust that is not based on acquaintance but on assumed common values and goals (for details on trust in networked production cf. Mezgar 2003). The problem here with regard to collaboratively produced Internet security is that swift trust might be inappropriate for sharing secretive information and for globally distributed collaborative communities beyond the scope of intra-firm departments or inter-firm collaboration amongst a few corporations.

1.3.2 Practical problem

The characteristics of current Internet security problems, their global distribution, and their requirements for operational speed, arguably demand transnational, rapid operational solutions based on many informational sources and a wide range of global knowledge. Therefore, a distributed approach appears to be a viable path for

to demand for it, will be used most efficiently" (Benkler 2002, p. 437). More extensive discussions on policing and security provisioning in (Crawford 2006, p. 466).

Internet security governance. The question whether and to what extent modes of peer production can be applied for Internet security is of interest for the design of the future organisational Internet governance landscape.

Policy makers have to decide about the future outlook of this regulatory field, and balance the degree to which states want to regulate and play an operational role in Internet security production. To achieve best regulatory results, the full regulatory and policy arsenal should be known and well researched. Just as governments have discovered the utility of open source software for public information systems, they might want to re-evaluate the peer production approach to achieve Internet security. Businesses, the financial industry and ISPs might want to re-evaluate their collaborative strategies with regard to Internet security governance, just as businesses have agreed before on sharing the code of commodity software. Even the ICT security industry or more precisely, anti-virus companies, have adopted new sharing strategies. While in earlier years AV companies treated their libraries of collected malware as proprietary, scarce assets, they eventually started sharing malware samples presumably as a response to their initial inability to stop the large virus floods in the early 2000s.

Applying modes of peer production for Internet security raises the urgent question of how to deal with the need to keep certain kinds of information secret and to deny certain actors access to the informational resources of distributed collaboration. Security in modern nation states has traditionally been in the hands of dedicated authorities, of which a distinguishing feature is the ubiquitous application of secrecy and restricted outbound flow of information. In settings in which insecurity is caused by actor-based threats, secrecy is a technique to deny adversarial actors a tactical advantage based on an informational edge. This clash between the openness in peer production and exclusion and secrecy in security politics raises the question of whether peer production models are applicable at all in the sphere of Internet security and if so, how actors can reap the advantages of peer and open source production while maintaining secrecy and avoiding the detrimental effects of disclosure and openness.

The characteristics of Internet security and the general advantages of peer production suggest that this production model is likely to be applied in the domain of Internet security beyond existing particular examples. This research aims at providing insights into existing forms of Internet security production and questions the role and existence of open forms of collaboration therein.

1.3.3 Peer-produced security

The idea of applying the model of peer production to counter some of the Internet security challenges is as intriguing and consequential as it is counter-historical, strange, and naive. It is intriguing, because the approach could ameliorate the long-lasting democratic deficiencies of traditional security organisations, by putting security back into the hand of the people. It is naïve, since the peer production of security would equal a transfer of power away from traditional security organisations to some potentially uncontrolled self-governed networks or communities. It is naïve too as foreign policy elites, at least in the US, have seen the Internet as a strategic resource for national power in International politics (Schmidt 2004). The recent leaks on the NSA/GCHQ Internet surveillance programs only underline this argument. Furthermore, theorists of the liberal democratic state would probably either laugh or cringe at the idea of incumbent elites voluntarily accepting a partial power vacuum or permanent transfer to another player in a new security domain created by the rise of the Internet.

And yet, the very rise of the concept of security governance (Krahmann 2010) indicates that states play a different, arguably lesser role in the security domain than in previous decades. The neoliberal turn in international security has nurtured the rise of new private security companies. It is far from clear whether or not the empowerment of the individual and of self-organized groups has, does, could or will leave its marks in the organisation of Internet security provisioning. Irrespective of such theoretical ruminating or empirical observations of an increased role of the state in dealing with the broad Internet security complex, there is at the same time the somewhat opposite phenomenon of a vivid self-organised network of Internet security experts, freely sharing information, knowledge and resources to effectively address Internet security problems. This kind of collaboration appears to have some of the characteristics of networked governance, of peer production or open source production.

Phishtank.org is a collaborative, open effort to gather and provide information on the latest phishing attacks. *Whois* is a publicly accessible database linked to domain name registries that contains contact information for every registrant. In the eyes of many Internet security pundits, its openness makes it an indispensable resource for countering and mitigating current Internet security problems. Distinct Internet incidents like the advent of a new botnet or large-scale attacks on a single country's Internet resources are addressed by expert communities, by bottom-up networks of individual actors and groups. A group of "unsung heroes save[d the] net from chaos" and re-established its functionality briefly after a substantial part of the Internet traffic was routed to unintended destinations.[5] In 2011, the free malwr.com service was launched. Security experts can send files potentially containing malicious code to the service and receive an automated analysis of the capabilities and features of the submitted malware.

More significant though for this research project is the collaboration among Internet security experts. They share intelligence information on recent attacks, data on the current health status of the Internet, information about current malware, viruses, phishing attacks and botnet infrastructures; they create, host, and use tools for operational and intelligence work. Thus, by sharing information, knowledge and tools, they help to improve the overall technical security of the Internet. At first sight, these collaborative efforts have a strong resemblance to the way in which open source software is produced. Hence, can Internet security be produced in the same way as its core technical foundations and software? There seem to be differences, as access to these incident response networks seems to be restricted to persons employed by certain organisations.

1.4 Research plan

The discussions in the previous sections have carved out the research problem that this study seeks to address: the unknown limits of peer production, its unknown applicability to the domain of security production, and the unknown organisational results when an open-source-prone constituency is exposed to the needs for secrecy. This study aims at untangling these research problems, guided by the following research questions.

The first and main question is: *Can the handling of a particular Internet security incident be classified as peer production?* The preceding subsection 1.3.3 has already indicated that incident response endeavours resemble classic open source software projects to some extent. The empirical sections of this study therefore dive into the response activities of two large-scale incidents and the responses to them to see where they correspond to, and where they deviate from, defined characteristics of peer production. The answer will increase our knowledge about the relevance of peer production in Internet security operations.

[5] Jonathan Zittrain quoted in Fildes 2009. Fildes reports on a talk by Zittrain, in which the latter mentions the role the NANOG community in managing the frequently mentioned Youtube-Pakistan-BGP incident (cf. Singel 2008).

Additional research questions seek to deepen our understanding of the role of secrecy and one of its main drivers, the antagonists, for the handling of Internet security incidents. When question 1 is answered with a clear-cut "yes", research question 2 highlights the seemingly contradictory peer-production/openness and security/secrecy couple. Question 2a asks: *If secrecy is of any importance for security provisioning, which techniques are used to retain both secrecy and openness?* Question 2b has a similar objective, asking: *If the antagonist is of any importance for security provisioning, which techniques are used to exclude the antagonists?* This question aims at understanding how a collaborative platform could retain a substantial degree of openness while being the potential target of the antagonist's retaliation.

The next research question applies when question 1 is answered with a "no", i.e., when the handling of an incident cannot be labelled as predominantly peer production. In this case it is of interest whether the response endeavour has some of the characteristics of peer production. Question 3a asks: *Which elements of the mode of peer production are used for security provisioning and for what purpose?* The answer to this question should provide more detailed information about the organisation of incident response activities. Question 3b is intended to yield answers about the hindrances of peer production: *What seems to be the primary cause(s) of the non-application of the mode of peer production?*

This study has several objectives. Some of them are directly related to the knowledge that should be yielded by answering these questions; others are the effect of the steps necessary to get to the position of being able to answer them in the first place. Answering the research questions apparently contributes to both the theoretical and practical problem regarding whether peer production actually is a viable organisational approach for incident response activities. Generalising on these findings, more could potentially be said about a possible role for peer production in security provisioning in general, be it related to the Internet or not. This study thereby contributes to the emerging body of literature on networked governance, peer production, and other forms of social production. Finally, this study researches the intricacies of the empirics of Internet security production.

The topic of this study is of scientific relevance for a number of research domains. Apparently, the existing literature on peer production provides little detail on the limits of this organisational approach, the potential hindrances to its applicability in domains other than software production. We know little about the organisational forms that develop when distributed, collaborative, bottom-up collaboration is used in an environment that appears to be friendly and hostile to peer production at the same time. More clarity on the limits and possibilities of peer production might lead to better practical judgement about investments in future research on production and governance innovation and in implementations thereof. The societal and political relevance is in another league. Weber has already stressed the potential international effects of a wide-spread adoption of the open source method. "For international politics, the demonstration of large-scale non-hierarchical cooperation is always important; when it emerges in a leading economic sector that disregards national boundaries, the implication are potentially large" (Weber 2004, p. 227). If that large-scale non-hierarchical cooperation would be applicable in the domain of security, it could possibly revolutionize the relations between citizens, traditional security institutions, networks, and states on both national and international levels.

The lack of academic rigour in this openness-debate has resulted in a battle of arguments that are too often based on anecdotes: hyped on the one hand by evangelising Internet intellectuals like Clay Shirky with claims that the Internet-facilitated collaborative turn revolutionises everything (2005, 2008); and dismissed on the other by snarky critics like Evgeny Morozov, who criticises his über-techno-optimist opponents as being "so open it hurts" (2013, ch. 3). Enriching our knowledge about the limits and possibilities of openness is of both scientific and practical relevance. The scientific elements have been stressed earlier in this chapter. The practical dimension results from the political discussions regarding the use of all facets of openness in all sorts of policy and polity dimensions, with open data and open government as only two examples.

This research also seeks to enrich the debate around existing Internet security problems and governance approaches to overcome them. Detailed academic accounts as to how Internet security is actually produced after large-scale security incidents have by and large been missing. Whether the narratives provided in this study will eventually be of any practical political relevance is arguable. They will, however, surely enrich the emerging historiographical literature on past Internet security incidents and enlighten what has so far been the black-box Internet security production. Furthermore, the study will provide insights into existing, but rarely noticed security institutions. Knowledge about them might be useful in debates on future organisational security architectures for the Internet and in debates into public-private partnerships on Internet security.

This study is organised as follows. The following chapter, Theoretical Foundations, dives into the theories, models, concepts, and previous research on the topic of this study. The literatures on Internet security, security governance and provisioning, and social and peer production are analysed. The goal of this chapter is to provide a model of peer production of Internet security, which is necessary for the further conduct of this study.

The chapter Research Design explains the design decisions made for the conduct of this research endeavour. This study features a case study approach, used to explore how Internet security production works and the role that the mode of peer production and its defining characteristics play in the overall organisational approach. The chapter develops an analytical model that is used in subsequent chapters to identify these characteristics in the response activities. A few remarks on the methods of data collection and analysis conclude this chapter.

In the chapter Endangering the Internet, the narratives of the attacks on Estonian Internet infrastructure in 2007 and the Conficker botnet are offered. In late April until mid-May 2007, a significant proportion of Estonian ICT systems providing Internet services were attacked, massively hampered, or even interrupted by various malicious methods. Roughly two years later, one of the largest of its kind, the Conficker botnet, peaked in activity. It infected millions of machines worldwide, and handed a massive network of potentially malevolent bots to the controlling masterminds behind the scenes. The ensuing chapter Producing Internet Security then tells the story of the responses to these attacks, and how the functionality of the Internet was re-established. These two chapters prepare the scene for answering the actual research questions.

The chapter Social Dimension of Internet Security applies the models and methods developed in preceding chapters. The response activities are analysed through the lenses of the operationalised model of peer production. The chapters provide detailed depictions of the decentralized nature of the response, the rules of sharing of resources, the role of proprietarity, markets, and hierarchies in the response endeavours of the two cases.

The chapter Limits of Openness analyses the role of secrecy and other factors limiting the applicability of peer production in the two cases. In addition, it elucidates the reasons why certain elements of peer production could be observed in the cases while others were absent.

Eventually, the Conclusion chapter sums up the findings of this study and elaborates upon its implications for the theory and practice of peer production and Internet security.

2 Theoretical Foundations

The introductory chapter briefly touched on the seemingly aligned trends, namely the increasing diversification and denationalisation of security governance on the one hand and the Internet-empowerment of distributed, self-organised, bottom-up collaboration on the other The peer production of Internet security would be a perpetuation of these trends. And yet, the usual secrecy that appears to come by default with security governance and operations runs across the definitional characteristics of peer production or other forms of open social production. Secrecy and openness are clearly opposing theoretical concepts. The limits of the applicability of peer production in other societal domains than software needs better understanding. Likewise, it is unknown which forms of social production might emerge, when ideas of distributed, self-organised, bottom-up collaboration are applied in the domain of security provisioning.

This chapter identifies and discusses the theoretical foundations and core concepts that are relevant for researching the relationship between peer production and security governance, or, more generally, the social organisation of security production. This study follows a multi-disciplinary approach, using literature on open source and peer production, the economics of Internet security, security and policing, and Internet governance.

Section 2.1 reviews the broad literature on new forms of collaboration. It analyses peer and open source production, extracts their core characteristics and their premises, and discusses the viability and applicability of social organisation in areas besides open source software production. The relative advantages and disadvantages of peer production over other modes of production and governance are also scrutinised. Furthermore, the section discusses weaknesses and limitations of the existing literature on social production.

Section 2.2 examines the wide field of Internet security governance, which overlaps with other complex fields such as Internet governance, security governance, and ICT security. It shows the increasing diversification of governance, security governance and the provisioning of public services. However, the full range of collaborative and governance forms presented in the first section have not been integrated into the literature on Internet security governance despite indications of their empirical existence.

The third section analyses how secrecy and openness relate to peer production and security. Therefore, the different meanings and social functions of secrecy are presented, and the ambiguous relation to technical security portrayed. The section primarily attempts to theoretically anticipate the possible impact of secrecy on the feasibility of peer production.

The fourth section finally integrates these previous streams of literature to overcome their respective deficiencies, our lack of insight into the limits of peer production and the lack of empirical analyses of actual production of Internet security. This research assumes that the peer production of Internet security is a distinct, available form or method of Internet security governance. This section presents a model of peer production of Internet security and its viability and hindrances.

2.1 Social production and the collaborative turn

With the advent of new technologies that are applicable to a wide range of possible fields of application, any existing societal institution potentially comes under pressure to adapt to new technologies or even be replaced by entirely new institutions. Information technology and the Internet have already resulted in a series of organisational and political changes (cp. section 1.2). A particularly noteworthy trend has been the rise of social production and peer production, a subtype of the former, referring to distributed, self-organised, and bottom-up-governed production.

The rise of peer production is part of a trend towards forms of governance and production that no longer rely on markets and firms, but increasingly on networks and hybrids that rely upon networks. Pervasive information technology and the Internet have facilitated new forms of geographically distributed communication, cooperation, and collaboration. These changes have appeared in probably all societal domains, and have been analysed in a wide range of academic disciplines, such as public policy, international relation, organisational theory, sociology, or economics.

Social production relies on social relations among contributors, not on economic incentives or hierarchical orders. It comes in different types and is labelled with different concepts such as distributed collaboration, open source method, or peer production. These forms vary, among other criteria, in their openness, distributiveness, and socialness. For peer production in its strict commons-based variant, the product is openly accessible, and access to it not restricted. For certain types of crowd-sourced production, however, property rights for the product remain with the platform-owner. Similarly, information and input resources necessary for production can be proprietary and therefore undermine the feasibility of a production model based on merely social incentives. This section discusses various forms of social production and similar types of collaboration.

2.1.1 Distributed collaboration

Distributed collaboration has become a major topic in economic, organisational and governance literatures, which study social organisations. New information technologies and forms of their usage have led to the rise of previously unknown distributed forms of collaboration. Distributed collaboration describes close collaboration among geographically and organisationally distributed persons or teams facilitated by modern information and communication technology. Open source software development and community-based content production are among the most prominent domains of application of this collaborative model (Shirky 2005, 2008; von Hippel 2002; Von Hippel & Von Krogh, 2003; Lakhani & Von Hippel, 2003; Cooper 2006; van Wendel de Joode, de Bruijn, & van Eeten, 2003). Distributed collaboration can include classic inter-firm collaboration as well as user-firm networks or open source-like networks of voluntary producers. Distributed collaboration among producers is supplemented by enhanced cooperation among distributed commanders. "Distributed decision-making" or "collaborative command and control“, applied, e.g., in military and emergency management, reflects the distributed collaboration among leaders of different, distributed, and even independent teams (Chumer & Turoff, 2006; Trnka & Johansson, 2009).

Examples of the combination of open source (as a software attribute) with distributed inter-firm collaboration are so called "open source service networks". They describe international networks "of firms that collaborate in order to service customer software needs based on open source solutions" (Feller, Finnegan, Fitzgerald, & Hayes, 2008, p. 476). These networks rely on a number of governance techniques and social mechanisms to coordinate and safeguard the exchange of information: restriction of access to the network, a "macroculture" of shared assumptions and values, collective sanctions against violation of shared norms, and reputation (2008, pp. 479-480) . These open source service networks differ from open source projects as their goal is not to produce open source software and they hence do not apply an open-source mode of production. Other than in conventional inter-firm collaboration, they use social mechanisms like reputation instead of legal means such as contracts.

Another broad concept for geographically distributed collaboration has been suggested by a joint research project by the Oxford Internet Institute and McKinsey.[6] In so-called "distributed problem solving networks" or "collaborative network organisations," "peers" and "networked individuals" collaborate on "problem solving and co-creation of services and products beyond traditional organizational boundaries and geographical constraints."[7] In his "classification framework", Dutton distinguishes between three types of collaborative network organisations, "1.0 Sharing", "2.0 Contributing", and "3.0 Co-Creating" (Dutton 2008, pp. 216-219). Each of these types also "links" with four different forms of "management strategies to collaboration", namely architecture, openness, control, and modularisation. As an example, Dutton asserts that the need for access control in co-creation networks like Mozilla's Firefox is greater than in sharing networks such as Bugzilla (2008, pp. 224, 217). While this stream of publications has made extensive use of the concepts used in this study, they have added only little conceptual clarity and theoretic value to the pre-existing literature described in the subsequent subsections.

A third way to frame distributed collaboration has been brought forward by Adler and Heckscher (2006). Similar to the conceptualisations mentioned above, Adler and Heckscher's "collaborative communities" aim to encompass broader social realities than peer production projects and their specific definitional requirements. In some societal areas, the authors argue, communities have become "the dominant organizing principle", superseding markets and firms (2006, p. 16). This "new form of community" (2006, p. 13) contrasts with "older" forms of community that were either in the shadow of markets or in the shadow of hierarchies. A collaborative community is based on values ("participants coordinate their activity through their commitment to common, ultimate goals"), organisation ("social structures that support interdependent process management through formal and informal social

[6] The project website has been available at http://www.oii.ox.ac.uk/research/project.cfm?id=45, accessed in June 2010.

[7] Dutton 2008, p. 211. For further project papers cp. Den Besten, Loubser, & Dalle, 2008; Loubser 2008.

structures") and identity ("reliance on interactive social character and interdependent self-construals") (2006, pp. 16-17).

Adler and Heckscher address fundamental sociological questions about communities in contemporary societies. Their observations on the characteristics of collaborative communities might help explain some of the potential empirical results of this study. "Neither the traditional nor modern forms of community are adequate for groups that seek high levels of adaptiveness and complex interdependence. In such situations trust is particularly important, because people depend a great deal on others whose skills and expertise they cannot check.… Collaborative community forms when people work together to create shared value. This increasingly characterizes societies in which the generation of knowledge, often involving many specialists, has become central to economic production" (2006, p. 20). Firms are affected, too, by the challenges the creation of knowledge creates for businesses. "Corporations are finding that to produce the complex forms of knowledge increasingly needed for economic growth — bringing together the expertise of multiple specialists — they need to move beyond the informal links of the paternalist community" (2006, p. 31). This need has apparently led to the corporate support of distributed, peer-governed production networks. The question arises however, as to how these networks are or could be organised in the security domain.

Adler and Heckscher explain why distributed forms of collaboration work despite the fact that the actors within them have never met each other in person, and nor is there a hierarchical intermediary that would guarantee and enforce mutually expected behaviour. The glue holding individuals together in "collaborative communities" is "swift trust" (Adler & Heckscher, 2006; Osterloh & Rota, 2004), a mode of trust that is not based on acquaintance but on assumed common values and goals. But is "swift trust" sufficient when it comes to security issues? With regard to collaboratively produced Internet security, swift trust might indeed be inappropriate for sharing secretive information and for globally distributed collaborative communities beyond the scope of intra-firm departments or collaboration amongst a few well-acquainted corporations. This raises the question of how trust is generated in distributed, heterogeneous global security networks, and whether secrecy is an indispensable attribute of security production. The first question will be addressed later in this thesis in section 7.3 on Trust in the security community; the latter in section 2.3 on Secrecy vs. openness below in this chapter.

2.1.2 Open source production

The rise of GNU/Linux and Free/Open Source Software (FOSS) has attracted the attention of social sciences. Its status as a phenomenon worth studying is due not only to the sheer output and results of projects, which are at least partly conducted by volunteers dedicating their spare time to create public goods, but also because it has created new forms of organisation and governance techniques for distributed web-based collaborative endeavours. The resulting body of literature should thus pose a fruitful source of ideas and models to analyse distributed, collaborative production in the domain of Internet security.

Many aspects of FOSS relevant to social sciences have been studied. Riehle (2007) analysed how open software has changed the behaviour of actors in the software market. Cooper (2006) described the sources of economic advantage and how savings in digital production processes can be made on supply-side resources and transaction costs and how demand-side values can be enhanced by applying open-source software and production modes, peer-to-peer and mesh technologies. Schweik, English, and Haire provide insight into the relevant factors that make open source collaboration projects successful, and how this form of collaboration can be successfully applied to non-software activities (Schweik, English, & Haire, 2008b, 2008a). The quantitative dimension of open source software has been analysed by Deshpande and Riehle (2008). Maxwell (2006) argues about the role of openness in open source software and its related property regime for the open-source innovation model. Riehle (2010) considers the use of the open-source model as an opportunity for firms to lower production costs and product prices, which would increase the overall market size for their products.[8] The consequences for openness of source code is discussed by Hoepman and Jacobs (2007), who conclude that in the long run openness of systems would make them more secure, while in the short run exposure would likely increase.

Given the challenges posed by Internet security risks and the technical and organisational innovations that are necessary to overcome them, the relationship between innovation and the open source method are of high interest. Open source software as a showcase for "open innovation" was analysed by West and Gallagher (2006). This mode of innovation "systematically encourag[es] and explor[es] a wide range of internal and external sources for innovation opportunities, consciously integrating that exploration with firm capabilities" (2006, p. 82) and transforms the management of intellectual property by firms. The open innovation model, as it

[8] For IBM's motives to engage with open source software cf. Samuelson 2006.

manifests in open source software development, has two key characteristics, namely collaborative development and shared rights to use that technology (2006, p. 91). Economic incentives for firms to participate in OSS development is to a) use open source projects as a means to pool their R&D resources,[9] and b) use existing open source projects as a foundation for commercial products and services. In a brief journal article, Mayer-Schönberger (2009) takes a different stance regarding the innovative potential of the open source method. He argues that disruptive technological innovation would likely be hindered, when a network or community is characterised by many-to-many connections. Dense networks and high connectedness would create "groupthink" and lead to incremental small-step-style innovation instead of what would be necessary to overcome spam and other challenges. "To enable innovations, especially non-incremental, discontinuous, and radical ones — which are needed, among other things, to launch successfully the next-generation Internet — may require unique policy intervention: reducing the social ties that link its coders." (2009) In contrast, Wang (2007) observed that the high-connectedness of its coders and contributors increases the likelihood that an open source project is successful.

Among the core issues of research on FOSS are questions about the nature and defining characteristics of the open source method, the sources for its success, and factors of its viability and sustainability. Osterloh (2004) identified intrinsic and extrinsic motivation and favourable governance mechanisms, which would not hinder the former, as prime factors for the success of open source projects. A set of motivational, situational, and institutional factors are prerequisites for the functionality of "virtual communities of practice". These communities are characterised by an absence of central authorities, and of privatisation of intellectual property rights, by loosely defined group borders and unclear resources (Osterloh, Rota, & Kuster, 2006). As to the motivational factor, Osterloh argues that actors need to be motivated by a mix of intrinsic and extrinsic factors. As to the situational factor, open source production is more likely to be successful when it is less efficient for independent suppliers to ensure private ownership of intellectual property. As to the organisational factors, volunteering, self-governance, participation and transparency of decision-making are supportive for open source systems. Finally, with regard to the institutional factors, license arrangements like copyleft and the support of open source community norms by commercial providers foster a successful application of modes of open source production (2006, pp. 23-27). David and Shapiro found that contributors to open source projects have a wide range of motivations. The degree of success of open source projects depends on their ability to

[9] The contributions of IBM, HP, and Sun to the Mozilla project served as prime examples.

rise and sustain motivations for actors to start and continue contributing to open source projects (David & Shapiro, 2008).

Factors that would make the *Success of Open Source* have however most thoroughly been studied by Weber (2004). The remainder of this section is devoted to those factors identified by Weber that make modes of open source production successful. Open source as "a way of organizing production" (Weber) — as opposed to open source as an attribute of software — is defined by some key characteristics: Everyone can participate and contribute to these projects, projects are set up and run on a voluntary basis, contributors share what is called the 'open-source ideology' and projects are organised around a certain property regime (2004, p. 268).

Valuable for the analysis of power-saturated security governance settings, Weber has discussed the role and locus of power in different kinds of networks. Differentiating between three network types (open source, platform, subcontracting), Weber sees power in an open-source network residing with those inventing and dominating the ideas and values of the projects. The ordering principle would be a voluntary, charismatic meritocracy, presumably with Linus Torvalds as the incarnation of that archetype (2004, pp. 257-259). Apparently, actual implementations of the open source ideal do come with some degrees of internal organizational hierarchy and do not resemble anarchic or power-free venues. The role of authority has been discussed in a number of further studies (Dafermos 2012; Konieczny 2010; Kostakis 2010; Viégas, Wattenberg, & McKeon, 2007; Loubser 2008).

Weber argues that the open source mode of production is a specialisation of "distributed innovation". This form of innovation is based on the four principles of experiment empowerment, mechanisms to identify relevant information, mechanisms to recombine information, and a governance system supporting such an innovation approach (Weber 2004, pp. 234-235). The distributed open source model empowers potential contributors to play with the given, freely accessible resources of an open source project, and recombine ideas and previous results to come up with new innovative ideas. The absence of central-decision making "in the sense that no one is telling anyone what to do or what not to do … is the essence of distributed innovation" (2004, p. 233). The innovation aspect of the source production hence relies on an appropriate structuring of information.

Most relevant for this study are Weber's contributions to a model of open source production viability. Weber has not developed an empirically tested theory of open source production, but he has formulated a convincing set of assumptions on factors influencing the feasibility of open source adoption in other domains than software development. He stresses that much of his thinking about the effectiveness and viability of the open source process in general are “expansive hypotheses”. But “they are broadly indicative of the kinds of conditions that matter” (2004, p. 266).

Based on and extending his empirical analysis, Weber has distilled a list of factors that likely allow “the open source process … to work effectively” (2004, p. 271). These factors are crucial for the viability of the open source production process in a particular domain of knowledge and for a particular good. The factors can be divided into attributes of tasks and production processes, and attributes of agents involved. Regarding the tasks and the general production process, effectiveness is increased by the following six factors: First, the knowledge necessary for the production process is *“not proprietary or locked-up”* and hence allows for “disaggregated contributions”. Second, the problem to be addressed is of a certain degree of *general importance*, e.g., a certain threshold in the number of beneficiaries is exceeded. Third, quality of production is ensured by mechanisms such as “widespread *peer attention and review*”, which manages to bring the number of errors introduced below the number of errors corrected. Fourth, “strong positive *network effects”*, partly supported by the aforementioned criteria, increase the effectiveness of open source as process. Fifth, given the voluntariness, open source processes tend to be more effective if they can be *initiated by “an individual or a small group* [that] can take the lead”. Last but not least, the production process is framed and supported by a “*voluntary community*” (2004, pp. 272; italics added).

Regarding the agents involved in open source production and the effect of their attributes on the effectiveness of open source production, actors need to be able *judge the viability* of an open source project. By contributing to an open source project, agents make a bet that joint efforts will result in a “*joint good*”. In addition, agents will expect *payoffs* for their contributions by “gain[ing] personally valuable knowledge” of positive normative and ethical valence (2004, pp. 272; italics added). Not only individuals can decide to participate in distributed innovation endeavors, but also organizations. The *feasibility* of distributed innovation is, according to Weber, increased by three further factors: the *problem resides within an organisation*; *more actors are likely to face the same or a similar problem;* and the solution to the problem does *not represent a competitive advantage* over competing actors (2004, pp. 265-266).

Weber builds his arguments partly on transaction costs economics, but he stresses that factors other than efficiency are of importance, too. He doubts that transaction costs are the main driver for organisations which are considering the option of open sourcing some of their activities. “Because no one can measure these activities in advance (and often not even after the fact), the decision becomes an experiment in organisational innovation.“ (2004, p. 266) A counterargument here could be that an organization's decision to open source certain activities would not be sustainable if it would not result in superior innovations, higher efficiency, and lower transaction costs in the long run.

Weber correctly identifies "two factors that transaction cost theory does not emphasize", including "the relationship between problem-solving innovation and the *location of tacit information*" and "the importance of *principles and values*, in contrast to efficiency" (2004, p. 267). Therefore, an organisation's or, more neutrally, an actor's decision to go open source is not merely an "efficiency choice around distributed innovation" (2004, p. 265). Much depends on the structuring and organization of knowledge in that domain. The distribution of *tacit knowledge* across several independent actors, the *necessity of such knowledge* for problem-solving, the *culture of both sharing knowledge* and relying on that shared knowledge are different in some societal domains such as medicine than in others, like software (2004, pp. 268-270). Sharing and using knowledge reflects *shared principles and values* among collaborators at least regarding how knowledge should be distributed in a particular domain or community (2004, p. 267). Weber discusses the idea of distributed problem solving among physicians and genomic researchers. He highlights how a barely developed sharing culture in the communities of both physicians and genomics scientists has stymied distributed collaboration in these two branches of the medical domain. Another stumbling block for the open source method in genomics is governmental regulation. This is despite the "structure of medical knowledge", which is widely distributed among countless experts. A superficial glance at traditional security governance practices suffices to hypothesize that the domain of security is packed with regulatory and cultural hurdles for the open source method. On the other hand the structure of security-relevant knowledge is widely distributed and the "key issue [of] organization of knowledge sharing" (2004, p. 268) has been solved in favour of the open source model in the information technology domain.

To sum up, a wide range of factors act as determinants of an institutional design. Whether the list of these factors is sufficient, the overall relevance of a respective factor on the viability of open source production is not clear. Neither is it apparent which institutional design will evolve or actors will chose when a few of the aforementioned factors are detrimental for the viability of the open source method in a particular context, while the majority of them support it. It is not clear which factor is *nice-to-have* and which ones are *sine qua non*. Until then, the voluntary application of the open source method as an organising principle in other domains than software is "experimenting" and "requires learning by doing" (2004, p. 268). Weber's decade-old question is still valid: "What happens at the interface, between networks and hierarchies, where they meet?" (2004, p. 262)

2.1.3 Peer production

Widely received in social sciences has been the idea of peer production, a term popularized by scholar of law and economics, Yochai Benkler, to explain collaborative arrangements for the development of free and open source software and Wikipedia (Benkler 2006). According to Benkler, open source software represented "a new modality of organizing production". To describe this new reality, Benkler uses the term "commons-based peer production", which is characterised as "radically decentralized, collaborative, and nonproprietary; based on sharing resources and outputs among widely distributed, loosely connected individuals who cooperate with each other without relying on either market signals or managerial commands" (2006, p. 60). The newness of that modality though is not, that "agents operat[e] on a decentralized, nonproprietary model", but the "increase in the importance and the centrality of information produced in this way" (2006, p. 63). Benkler argues that social production had become exponentially more feasible and widespread with the rise of a ubiquitous information infrastructure, which has fundamentally altered the economics of distributed, large scale information production and dissemination. Once expensive physical machines, distribution systems, and coordinating overhead bureaucracies were required for such production. But cheap Internet-based communication and highly distributed ownership of relatively cheap ICT systems have eradicated costs to set up and, combined with social production, operate such collaborative efforts.

Until Internet-based collaboration emerged, the term *peer* was used in the context of information production mainly in academic circles in the form of *peer review*. The term peer is deeply rooted in English social and legal history. The Magna Carta, composed in the 12th century, prescribed that "Earls and barons shall not be amerced except through their peers…"[10] The concept of peers refers to 'equals' or persons of the same standing. One of the first documented uses of the term *peer production* though happened to be a letter to the editors of the scientific journal *Physics World* titled "End of peer production", in which the author criticised the state of the academic peer review system (Venables 1991). The term *peer* was then picked up by technology circles to describe an emerging technology that allowed equal nodes in a network of computer clients to communicate with each other without the need for relaying central servers. These peer-to-peer technologies became widely prominent with the rise of Napster and similar sharing networks. Eventually, the term *peer* was used in a paper by Eric S. Raymond, a seasoned free/open source software developer who had developed an interest in reflecting on

[10] The Magna Carta, http://www.constitution.org/eng/magnacar.htm

the community of his *peers*. His 1998 article *Cathedral and the Bazaar* was the first or at least the first influential analysis of the organisational principles and peculiarities of free/open source software development (Raymond 1998a). In the subsequent paper *Homesteading the Noosphere*, Raymond used *peer* to describe fellow developers in an open software project or community, e.g., in the form of "reputation among one's peers" (Raymond 1998b). The term then spilled over to traditional academic institutions. The Berkman Center for Internet & Society, which is now presided by Benkler, published a research paper by David Bollier entitled *The Power of Openness*, which argued that the "open source revolution" was driven by "alternative economics" based on a "gift culture" and "voluntary behaviour" by *peers* (Bollier 1999). Steven Weber also used the term *peer* in his paper on the *Political Economy of Open Source*, but again only to describe fellow developers in software communities (Weber 2000).

It was Benkler who coined the term *peer production*, or more precisely *commons-based peer production,* to describe that new, emerging type of cooperation, production, and community. Benkler had studied the role of commons in information policies during the 1990s and published a few articles on this topic (Benkler 1998, 1999), when he brought his focus on new modes of collaboration facilitated by the Internet.

In *Coase's Penguin*, first presented at a conference in 2001 (Benkler 2001), Benkler analyses the economic foundations of peer production and argues that its viability is based on lower transaction costs (2002). "Peer production emerges, as firms do in Coase's analysis, because it can have lower information opportunity costs under certain technological and economic conditions" (2002, p. 374). Benkler sees the peer production model at work for several types of informational products and services, ranging from the production of content (the web itself as sort of an encyclopaedia, Wikipedia, and older projects like NASA Clickworkers, Everything2.com, Kuro5hin), relevance and accreditation (Amazon's user reviews, Google's PageRank, Slashdot) to value-added distribution (2002, pp. 381-399). In the two main analytical section of his paper, Benkler tries to answer the questions as to why peer production has emerged as a form of production next to firms and markets; and why contributors voluntarily contribute to such informational commons despite their often lamented, alleged tragedy.

Benkler's theoretical point of departure is transaction costs economics, and he basically applies Coase's way of reasoning the existence of the firm (Coase 1937) to peer production as an "important mode of information production". Coase argued in a rather non-Marxist way that firms exist because the transaction costs for coordinating agents, efforts, and resources within a firm by hierarchical command would be lower for some goods than by buying them on the market with its price system which however required some product specificity. Benkler argues that a peer production system has advantages over firms and markets in uncertainty reduction and "allocation efficiencies gained from the absence of property" (Benkler 2002, p. 406). Briefly worded, peer production has information gains and allocation gains over markets and firm-based hierarchies. Benkler carves out the information gain argument by modelling modes of production — i.e., firm-based hierarchies, markets, peer production — as "information-processing systems" (2002, p. 408). Knowledge-based work is based on individuals. However, "specification and pricing of all aspects of individual effort as they change in small increments over the span of an individual's full day, let alone a month, is impossible" (2002, p. 409). Agents, i.e., individuals, would be much better at identifying where they are needed and where their talents are of use. "The widely distributed model of information production will better identify who is the best person to produce a specific component of a project, all abilities and availability to work on the specific module within a specific time frame considered" (2002, p. 414). A second economic advantage next to these information gains are "potential allocation gains enabled by the large sets of resources, agents, and projects available to peer production". Benkler again uses a model to drive forward his argument. Based on it, he argues that the "productivity of a set of agents and a set of resources will increase when the size of the sets increases toward completely unbounded availability of all agents to all resources for all projects" (2002, p. 416). In a world of peer production, any of the countless individuals potentially willing to voluntarily contribute to any project can decide for which of the endless projects she is best suited and most needed.

The second question of the paper asks what motivates persons to voluntarily contribute to a commons system. Benkler argues based on a review of various types of literature and again some logical modelling that contributors are driven by the pleasure of creation and possibly also by "indirect appropriation mechanisms" such as reputation gains, consulting contracts, or a broadening of their skill set with potential effects on the job market in the future (2002, p. 424). Benkler's first argument is that under some conditions "peer production processes will better motivate human effort than market-based enterprises" (2002, p. 426).[11] In order for

[11] Benkler's formula for this is: $C_m > V > C_{sp}$ and $H + SP - C_{sp} > 0$ (H=intrinsic hedonic rewards, SP= social-psychological r., M= monetary r.; $C_{m/sp}$ ="costs for defining and making M and SP available to the agent"; V = "marginal value of an agent's action"). It means that the conditions for an agent to contribute voluntarily are that 1) the marginal value of an agent's action is a) higher than the transaction costs of defining and making monetary rewards available to an agent and b) lower than the transaction costs of defining and making social-psychological rewards available to an agent; and 2) the sum

Footnote continued on the next page.

peer production projects to attract high numbers of voluntary contributors, the project's workload needs to be split into several working modules with varying, heterogeneous granularity in order to attract contributors with varying motivations (2002, pp. 434-436). Furthermore, Benkler identifies a second set of risks for the viability of peer production. A failure of the integration process, in which all these tiny contributions are assembled to a coherent outcome, would demotivate contributors (2002, pp. 436-44).

In *Sharing nicely*, Benkler applies the idea of non-market and non-hierarchical based cooperation to the "domain of sharing rival material resources in the production of both rival and nonrival goods and services" (Benkler 2004b, p. 276). The article sounds out the "feasibility space for social sharing" of "mid-grained" or shareable goods, which are in "relatively widespread private ownership" and "systematically exhibit slack capacity" (2004b, pp. 276-277). Unlike markets and hierarchies, social sharing does not "require crisp specification of behaviours and outcomes" (2004b, p. 277). The first section presents two cases which exhibit social sharing of mid-grained goods or "sharable goods", namely carpooling and instances of distributed computing such as SETI@home. Benkler distinguishes "shareable goods" from *club goods* and *common-pool resources.* The latter two by definition are not "individually owned goods". The key characteristic of these shareable goods is that they "have excess capacity and are available for sharing" (2004b, p. 296). For Benkler, the form of production that is used for a particular good depends on a rather simple formula: "The combined effect of the motivational effects and the transaction costs of each system will determine, for any given good or use, whether it will most efficiently be provisioned through the price system, a firm, a bureaucracy, or a social sharing and exchange system" (2004b, p. 306). The transaction costs of social exchange differ from market exchange. Both require substantial set-up costs, but for social exchange, marginal transaction costs are lower as it only requires a lower "degree of precise information about the content of actions, goods, and obligations, and [less] precision of monitoring and enforcement on a per-transaction basis“ (2004b, p. 317). The second factor, motivation, is nurtured by social-psychological returns that are "neither fungible with money nor simply cumulative with it", but could on the contrary be diminished monetary rewards (2004b, pp. 326, 328). In the last main section of the article, Benkler argues that "sharing is a common and underappreciated modality" (2004b, p. 332) of production.

of hedonistic/intrinsic rewards and social-psychological reward are higher than the transaction costs of defining and making monetary rewards available to an agent (2002, pp. 426-431).

These two articles, *Coase's penguin* and *Sharing nicely,* make the basis and core of Benkler's thinking on peer production and social production, which culminated in the book *Wealth of Networks*. While the book certainly popularized his ideas and offered a more intuitive, narrative approach to the topic, the theoretical chapters strongly leaned on his previous articles on peer production and social sharing with their dense modelling and theorising approach. The book's theme again is that a "new model of production has taken root; one that should not be there." (2006, p. 59) This new phenomenon of "effective, large-scale cooperative efforts — peer production of information, knowledge, and culture" — had been applied in numerous projects (2006, p. 5). The explanations for the emergence of this phenomenon are supplemented by a normatively driven exploration of the possible impacts and opportunities of this new mode of production on individual, political, and cultural freedom, and on justice and development.

In the recent *Practical anarchism*, Benkler evaluates the overall relevance of the "peer model" and "voluntary open collaboration" (Benkler 2013, p. 213). He argues that peer production supplements existing markets and state-based provisioning systems; it does not replace them (2013, p. 245). Empirical instances of peer production would not need "a perfect or complete solution" to build "a degree of freedom into the world" (2013, p. 246).

A number of scholars have picked up Benkler's thinking about peer production and created complementary studies or modified the theoretical and normative basis of peer production. The paper *The accountable net* by Johnson & Crawford has arguably been the first to combine Internet governance and peer production. The authors argue that the Internet would not need "a more centralized authority" (Johnson, Crawford, & Palfrey Jr, 2004, p. 20), nor would representative democracy be "the best available form" (2004, p. 32). Internet governance should be congruent, and flexible (2004, pp. 31-32). Instead, the technical possibilities allowed for "a more powerful form of decentralized decision-making" (2004, p. 3), in which "we", the nodes of the networks or the peers, "take that law [i.e., referring to Lawrence Lessig, software code; A.S.] into our own hands" and create a "decentralised order" on the Internet (2004, p. 33). The paper acknowledges that some form of governance is required because of problems such as spam, spyware and other security issues (2004, pp. 7-11).[12] Daniel Neilson develops a theoretical, mathematical

[12] Articles that have explicitly tried to combine peer production and Internet security are discussed further down in this chapter in section 2.4.1, which also includes a more thorough discussion of the security aspects in Johnson (2004).

model to theorise about modularity, organisational hierarchy, and project design in peer production endeavours (Neilson 2009).

The debate on peer production has repeatedly been enriched by activist-intellectuals like Michel Bauwens. The former businessman and entrepreneur has left Belgium for Thailand where he has set up and heads the Foundation for P2P Alternatives. Since 2012, he also co-edits and publishes the web-based open-access *Journal of Peer Production*. Bauwens apparently prefers the term peer-to-peer or P2P over peer production, but in fact uses them interchangeably. In *Political economy of peer production*, Bauwens aims at developing "a conceptual framework ('P2P theory')" (Bauwens 2005). His conceptualisation of peer production or, in his words, "P2P processes", is also different and refers only to "those processes that aim to increase the most widespread participation by equipotential participant". P2P processes combine "free cooperation of producers [with] … access to distributed capital" with self- or peer-governance and open access to produced goods, ensured by new property regimes. He sees his P2P theory as "an attempt to create a radical understanding that a new kind of society, based on the centrality of the Commons, and within a reformed market and state, is in the realm of human possibility". Bauwens' explicit normative and policy goals apparently differ from Benkler's more moderate stance. While Benkler sees peer production as complementary to other modalities of production, Bauwens is more of an activist, and assumes that peer production would "overhaul our political economy in unprecedented ways" and help create a "powerful alternative to neoliberal dominance". The "nascent P2P movement … is fast becoming the equivalent of the socialist movement in the industrial age". In a more recent article, Bauwens introduced the idea of a *Partner state* with a reconceptualisation of the state's core tasks. A partner state would act as an enabler and facilitator of peer production activities (Bauwens 2012b, p. 39). Independent of one's political and ideational preferences, Bauwens remixing of states and networks is intellectually fascinating. It has "attracted" veteran network and netwar scholar David Ronfeldt to discuss and compare Bauwens ideas of P2P and the partner state.[13] So much for an overview of and introduction into the existing body of literature on peer production.

[13] David Ronfeldt, "Bauwens' Partner state", July 2011, http://twotheories.blogspot.com/2011/07/bauwens-partner-state-part-1-of-2-vis.html; David Ronfeldt, "In favor of "peer progressives": how, where, and why they're good for TIMN" (series of three blog articles), June 21, 2013, http://twotheories.blogspot.de/2013/06/in-favor-of-peer-progressives-how-where.html

2.1.4 Defining characteristics of peer production

When a new phenomenon emerges, it is not always exactly clear why it matters, what the defining characteristics are, which factors have facilitated the rise of this new phenomenon, and how it relates to existing, known phenomena. Different analysts perceive this new phenomenon from different perspectives, ideational and research backgrounds, normative stances, and policy agendas. And sometimes a new phenomenon comes in slightly varying flavours, which nevertheless can have profound consequences, which are differently gauged depending on the aforementioned perspectives. Peer production is one such phenomenon. A slight variation in the degree of control of one production resource can eliminate the normative promises of peer production as a sponsor of "virtue" (Benkler) or as an alternative to neoliberalism (Bauwens). I try to avoid such normative debates in this thesis. Conceptual clarity already is a sufficiently laborious topic; it requires focusing on the analytical dimension and finding appropriate, different terms for sufficiently different occurrences. This section first describes the ambiguities of existing conceptualisations of peer production. It then proposes a set of defining characteristics which will later be needed to identify peer production in Internet security.

The status quo of terms and conceptualisations of what one might call peer production is unsatisfactory. Different terms appear to be used to describe the same idea, the same term is used to describe different ideas. Benkler's recent *Practical anarchism* article is a poster-boy example of throwing a myriad of terms against an idea that lacks a thorough, clear-cut definition with identifiable criteria. The terms Benkler uses to describe the peer production model are: "peer mutualism", "peer-based efforts" (Benkler 2013, p. 244), "peer mutualism, or peer production" (244), "networked peer models" (214), "peer production practices and functioning mutualistic associations" (216), "cooperative systems" (216), "peer systems" (216), "voluntaristic associations…that do not depend on delegated legitimate force" (217), "voluntary open collaboration" (213), "peer networks that are voluntaristic, non-coercive, non-hierarchical, and yet productive" (217), "'working anarchy' or mutualism" (217), "mutualism, or practical anarchy, as an organizational form" (227), "commons-based peer production" (218), "voluntary associations of developers" (224), "peer production communities" (225), "large-scale, society-wide and economy-wide production that is based on a nonstate, non-proprietary model" (226), "commons-based peer production" (230), "decentralised, voluntaristic systems" (235). These terms combined signify some decisive characteristic of *the phenomenon*: peer-based, voluntaristic, non-coercive, mutualistic, non-hierarchical, commons-based, decentralised. It is not quite clear whether the terms above synonymously describe the same concept or describe very similar concepts. After all, these characteristics might not occur concurrently, one or more characteristics might be less pronounced. The question then is how should such an empirical or theoretical instance be called? Is a non- or semi-commons variant still part of *the phenomenon*, is it something different, or is it a subtype? Whether a variation of a characteristic requires a different term depends of the importance attached to that characteristic for *the phenomenon*.

While the use of dozens of synonyms for the same concept might be somewhat confusing, there is no major logical problem with it. An unclear relation between core concepts, on the other hand, especially when it borders on equivocation, is obfuscating or even misleading. Therefore it is necessary to clarify the relationship between the terms peer production, commons-based peer production, and social production.

Commons-based peer production can be characterised by "decentralized information gathering and exchange to reduce the uncertainty of participants", "very large aggregations of individuals independently" doing things, individuals who are "in search of opportunities to be creative in small or large increments", self-identification of tasks by these individuals, and diverse motivations (2002, p. 376). A similar definition by Benkler has been mentioned in earlier sections. Commons-based peer production is a "a new modality of organizing production", which is characterised as "radically decentralized, collaborative, and nonproprietary; based on sharing resources and outputs among widely distributed, loosely connected individuals who cooperate with each other without relying on either market signals or managerial commands" (2006, p. 60). The term *commons-based peer production* verbatim reflects two aspects of the production process, peer and commons. First, in *peer production*, production is performed by a group of peers, that is individuals acting on their own behalf and not coordinated by price system or managerial commands. In Benkler's words, peer production "refers to production systems that depend on individual action that is self-selected and decentralized, rather than hierarchically assigned" (2006, p. 62). Second, *commons-based* refers to the characteristics of the inputs to and results of the peer effort. Commons-based input resources and produced goods "are not built around the asymmetric exclusion typical of property" (2006, p. 62). Economic theory differentiates goods by a 2x2 matrix representing whether a good is rivalrous (or scarce or subtractable) and excludable. Non-excludable, non-rivalrous goods are called public goods, while non-excludable, rivalrous goods are called common-pool resources goods. Examples for public goods are sunsets and knowledge that is in the public domain; examples of common-pool resources are ocean fisheries, domain names or libraries (Hess & Ostrom, 2003, p. 120). In a more down-to-earth definition, *commons* is something that is "available to anyone who wishes to participate in the networked information environment outside of the market-based, proprietary framework" (Benkler 2006, p. 23). Commons in this sense means that a good is accessible to anyone and shareable within social transactions.

According to these definitions, there is a decisive difference between peer production and commons-based peer production. The commons, using Benkler's definition, is something "that is available to anyone who wishes to participate." C*ommons-based peer production* requires that the produced good is a commons, whereas plain *peer production* does not. The examples of peer production Benkler shares in his writings illustrate this. Take his *Coase's Penguin* article for example. The first main section titled "Peer production all around" mentions not only free and open source software projects and Wikipedia, but also NASA Clickworker, Google PageRank or Amazon consumers' reviews as examples. In *Wealth of Networks*, Skype was given as an example of social sharing of communication platforms for its implementation of P2P technologies (as opposed to P2P practices described by Bauwens) (2006, p. 421). Google PageRank is categorized as "peer production of ranking" (Benkler 2002, p. 392, 2006, pp. 76,171). PageRank is a publicly accessible numeric value expressing the computer-generated relevance of a webpage based on a proprietary algorithm owned by Google; the algorithm uses links, metadata and other user-generated values as input variables (Brin & Page, 1998). With the exception of search engine optimizers, users usually do not generate these values specifically as a contribution to the PageRank system. Following Benkler's logic, NSA's recently revealed systems, which trawl through any sort of user-generated information ranging from Internet webpages to backend data streams, are peer production, too. The only difference is that Google allows a glimpse into its computed data, the PageRank, while access to NSA generated data is restricted to a narrow intelligence community. Amazon's user-written reviews are different in several dimensions. First, contributors explicitly create texts for the purpose of being displayed as reviews under products offered by Amazon. Second, Amazon uses contributors' input in an unaltered way; reading users see exactly what other users have contributed — presumably minus pornographic texts and whatever else is deemed inappropriate and filtered away by Amazon. The examples of NASA Clickworker, Google PageRank and Amazon user reviews are similar in one important aspect: A corporation provides a proprietary platform that allows users to contribute some information that is henceforth used and presented — and sometimes even owned — by the platform owner at his liking. Such a definition of peer production substantially differs from Benkler's more rigid commons-based peer production. Bauwens on the other hand tries to avoid that hunky term by defining peer production similarly to how Benkler defined commons-based peer production, thereby avoiding the confusion as to whether peer production is used as the short version of commons-based peer production or as peer production *sui generis*. Bauwens list of characteristics of peer production is rather long and inclusive (Bauwens 2012a).

By calling the above-mentioned three cases — Clickworker, PageRank, Amazon reviews — an example for *peer production*, the term apparently includes what is often called *crowdsourcing*. In numerous instances however, Benkler uses the term *peer production* when he actually discusses an instance of *commons-based peer production*. Hence peer production at times appears to be the short form of the bulky term *commons-based peer production*. Crowdsourcing can be defined as "an online, distributed problem solving and production model already in use by for–profit organizations such as Threadless, iStockphoto, and InnoCentive." (Brabham 2008b, 2008a) But there are at least three different definitions of crowdsourcing used by scientists and journalists: In the first, narrow sense, crowdsourcing refers to a mode of production (or problem solving) in which a central firm harnesses selected contributions from individuals who respond to an open call for proposals (Howe 2006). Defined in this sense, crowdsourcing combines elements of Taylorism — by splitting up working packages into small chunks like in Mechanical Turk — with a tender technique often used by buyers of creative services: Service providers are invited to first show what their solution would look like and the orderer then decides whether the proposal, which entails a significant part of the entire solution, is worth some of his money. In this model of crowdsourcing, contributors compete, rather than collaborate or co-create (Woods 2009). In a second sense, crowdsourcing does not only call for individual proposals, but also includes self-organized collaboration amongst individuals of the crowd. Both this and the previous definitions assume the monetization and sometimes also the propertization of the contributions by the central organisational firm or actor. This second definition — Brabham has given an example thereof (see earlier in this paragraph) — assumes that collaborating groups can produce more innovative results than the ingenious virtuoso individual. The third is a catch-all definition for all kinds of "co-creation" akin to Benkler's implicit definition of *peer production* outlined earlier in this section (see the discussion of the section "Peer production all around" of his *Coase's Penguin* articles above on p. 33). For Kazman and Chen, *crowdsourcing* is "the popular term for commons-based peer production" (Kazman & Chen, 2009, p. 76). Kazman and Chen focus on the unifying element, "the crowds", which "co-create" in projects "from OSS to Wikipedia, Facebook, Amazon's Mechanical Turk, and many other community-based service systems" (2009, p. 77). Using these terms — crowdsourcing, peer production and commons-based peer production — as synonyms nullifies the opportunity to differentiate between production models that make different uses of hierarchy, control, collaboration, appropriation and monetary incentives.

Another term frequently used by Benkler is *social production*, especially in *The wealth of networks* as the subtitle *How social production transforms markets and freedom* indicates. Social production describes the third mode of production next to firms and markets, which they can either complement or replace (Benkler 2006, p. 117). "[M]arkets, firms, and social relations are three distinct transactional frameworks." (2006, p. 107) In his 2006 book, Benkler writes that social production includes "peer production of information, knowledge, and culture and sharing of material resources" (2006, p. 120). Social production — or in the longer form, "social production and exchange" (2004b, p. 336) — hence is the umbrella term for the two phenomena described in *Coase's penguin* and *Sharing nicely*, "commons-based peer production" and "social sharing of shareable goods."[14] But *social production* is also used in an even wider sense, pointing to any kind of transaction within the social framework and thereby ranging as far as Ostrom's famous Maine lobster-fishing gang (2004b, pp. 336-337).

Apparently, these at times overlapping and contradictory conceptualisations and taxonomies require some clarifications. The following conceptualisations and taxonomy might help to clearly label varieties of social production. *Social production* is the umbrella term, describing new non-market and non-hierarchy forms of production that have been facilitated by the rise of Internet and the decrease of transaction costs for distributed knowledge production and collaboration. The term stresses the *socialness* of a production process and is not based on market and hierarchical (as known in firms and other organisations) mechanisms. This conceptualisation excludes those types of crowdsourcing which are based on strong monetary incentives known from platforms like Amazon's Mechanical Turk. *Peer production* is reserved for that subset of social production that is based on a relatively egalitarian distribution of authority among the participating actors. *Commons-based peer production* is reserved for that subset of social production that meets the criteria of peers and commons; i.e., it is based on collaboration among equal peers, the absence of a central controlling governance unit; and it is based on the non-appropriation of produced goods. *Crowdsourcing* describes the outsourcing of tasks to an unknown crowd of contributors organised by a central organiser, which can

[14] This perception of Benkler's conceptualisation and taxonomy is reinforced by several phrases like "social production in general and peer production in particular" (2006, p. 123) or "why, and under what conditions, commons-based peer production, and social production more generally" (2006, p. 107).

be an individual, a group or an organisation. Contributions can be based on intrinsic and social incentives (*social crowdsourcing*), known, for example, from the co-creation of public intelligence after natural hazards. Contributions can also be based on monetary incentives like in iStockphoto (*market-based crowdsourcing*). Social crowdsourcing can be seen as a specialisation of social production. As crowdsourcing relies on a coordinating centre, it follows a decentral hub-and-spoke model. Regarding openness, social crowdsourcing can vary in accessibility of intermediary results, outcomes and the production platform.

The appropriate definition of peer production requires a more detailed look at the defining characteristics. Authors like Benkler or Bauwens have attached quite a number of characterising attributes or characteristics to commons-based peer production, ranging from peer-based, voluntaristic, non-coercive, mutualistic, non-hierarchical, commons-based, very decentralised, distributed, mass-based, large-scale, intrinsically motivated, creativity-driven, self-identifying with regard to tasks, forkable. The first research question of this research project asks whether peer production is actually applied in Internet security. Answering this question requires an operational set of defining characteristics. For this purpose, commons-based peer production is defined by the following three characteristics: *distributiveness*, *openness*, and *socialness*.

Table 2.1: Varieties of social production and their characteristics[15]

	Distributiveness	Openness	Socialness
Social production			★
Peer production	★		★
Commons-based peer production	★	★	★
Social crowdsourcing			★

Before detailing these characteristics, the underlying model of production needs to be presented briefly. (See Figure 2.1) The production model consists of input resources, the production process and the output products. Input resources are those goods that are required to be able to produce the actual output. The production process describes how the processing of input resources takes place; it also includes the production platform or facilities, tools, data, organisation, and intermediary products that are necessary to produce the actual output. The product or output

[15] The black star indicates a defining characteristic.

describes the intended result of the production process. All three basic elements can have several characteristics, which in combination characterise the entire production. For example, input resources can be unrestricted, shareable or not. The production process can be accessible for anyone or only for a restricted number of persons; it can be hierarchically organised or flat; contributors' participation can be driven by monetary incentives, hierarchical coercion or intrinsic motivations; organisationally, it can be widely distributed, decentralised or centralised; contributors can collaborate with or be separated from each other; contributors can act transparently amongst one another or only towards a central authority;[16] intermediary products can be openly accessible or proprietary. The resulting products can be proprietary or not. The white balloons in Figure 2.1 show the features that correspond to commons-based peer production in its purest form.

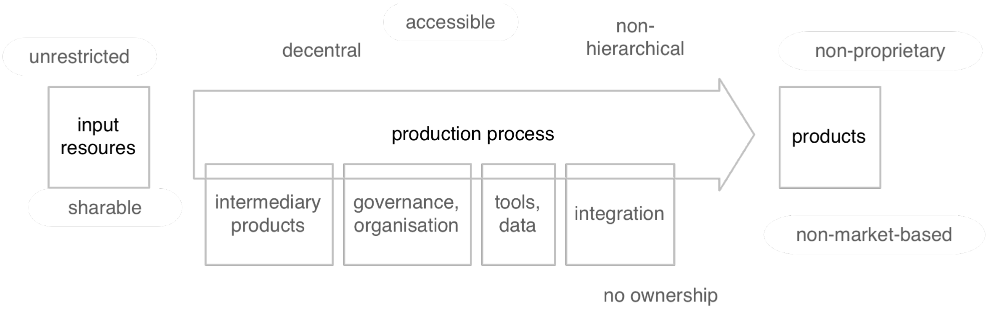

Figure 2.1: The underlying production model and characteristics of commons-based peer production

Distributiveness refers to the network topology of the contributing agents. Peer production is characterised by the absence of a central hub or even a limited number of hubs with substantial control. The characteristic of distributiveness also describes the model of peer governance, a distributed model of "self-regulation of peer groups, … creat[ing] new use value for the commons" (Bauwens 2012c). Peers can be defined as individual actors within a networked organizational form that have a relatively similar power status and ownership of assets in a particular environment. Hierarchies, if existent at all, are only ad-hoc or meritocratic. Crucial resources are also distributed among contributors, so that relevant explicit or tacit knowledge required for problem solving is not concentrated and therefore potentially scarce.

[16] Bauwens calls this holoptism and panoptism (Bauwens 2012a, 2005).

Openness describes the absence of restrictions to the accessibility of the production platform, i.e., the communicational space or physical facility, in which contributors meet to get access to necessary resources and tools, co-produce the intermediary or the output product, share information, or exchange ideas. Openness furthermore refers to the accessibility of intermediary goods and to internal, mutual transparency among collaborating peers about their activities and contributions (holoptism). Openness also describes the accessibility and modifiability of produced goods by any interested party, without having to utilise market- or hierarchy-based exchange frameworks. This aspect of openness, also called non-proprietarity, could be ensured by a commons-character of the produced good, and also allows for the forking of an informational product. An additional indicator is a shared sentiment among participants akin to the open source ideology (Weber 2004, p. 268).

Socialness refers to the social transaction framework on which the contributions of the participants happened. Contributors provide their tasks not based on monetary incentives or hierarchical pressure. Instead, they are driven by intrinsic and social-psychological incentives. Contributors act voluntarily and not based on market incentives.

Table 2.2: Details of defining characteristics of social production

Distributiveness	Distributed network-topology of contributors; absence of central hub or decentralised hubs Peer governance; hierarchies only ad-hoc or meritocratic
Openness	No or low access restrictions on production platform Accessibility of intermediary goods Internal transparency about activities and contributions Produced goods non-proprietary, accessible, reusable, adaptable outside market/hierarchy-exchange frameworks; forkable Open source-ideology
Socialness	Non-hierarchical Non-market-based Voluntary; intrinsic motivations

These criteria now allow for a better separation and a more transparent and logical definition and taxonomy of the varieties of social production. Following this taxonomy, peer production lies in between mere social production, which is characterised only by the socialness of the production process, and commons-based peer production, which adds strict requirements regarding the openness, i.e., the accessibility and non-proprietary nature of intermediary and final products. The difference between peer production and social production is that the former is more distributed and does not allow for centralisation by a super-node. Peer production can hence be seen as synonymous to open source production. They both span a wide range of phenomena from GPL-like property structures to Android with its countless anchors for Google to unilaterally steer and dominate the platform. Peer production relates to commons-based peer production as free software relates to open source software. With these steps towards conceptual clarity of social production, the question arises as to why this particular mode of production has arisen in the first place.

2.1.5 Feasibility of peer production

With theoretical discussions and empirical story-telling, Benkler convincingly argued that peer production is no one-hit wonder that only scored high in hippy software smitheries. Instead, peer production is a viable model for information production in general. Benkler has developed a model of peer production feasibility that is rooted in a mix of institutional economics and analyses of human motivations. His argument for the viability of peer production rests on four pillars: the characteristics of the networked information economy, lower information opportunity costs, the avoidance of potential hindrances for peer production, and norms in the form of regulation and cultural preferences.

Benkler locates his thinking in the domain of "information and culture". Therein, production required existing information as input resources; "human creativity" to create new information, knowledge, wisdom; and physical capital to edit, store, and transmit informational goods (Benkler 2002, p. 377). In his general argument Benkler paints a picture of the then only just emerging networked information economy, which fundamentally deviates from the pre-digital era of information production, the "industrial information economy" (2006, p. 32). The general argument is that the networked information economy has a few characteristics that facilitate peer production. Benkler mentions four characteristics that set the "networked information economy" apart from other economic subsystems (2002, p. 404). First, the rise of ICT, low costs of communication and ubiquity of production resources — a PC, an Internet connection, and some hosting space — had substantial lowered capital costs for information production. "[T]he physical machinery necessary to participate in information and cultural production is almost universally distributed" (2006, p. 105). Second, communication and information exchange among distributed contributors were cheap and efficient. Third, informational goods were non-rival. On the output side, they can be shared and used without being depleted. On the input side, the “primary raw materials…are public goods” (2006, p. 105) and can therefore be used at the liking of contributors. The “inputs necessary…are under the control of individual users” (2006, p. 105). What sounds similar at first, can also include scenarios in which input resources are not at all public goods, but have their access restricted to only the contributors. Fourth, creative talent was central for the production of informational goods and widely available (2002, p. 378). These arguments are pretty straightforward, and can be found in similar wordings in the writings of Internet-changes-everything evangelist-scholars like Clay Shirky (Shirky 2005, 2008).

The second line of reasoning goes deeply into the workings of information production processes. In his *Coase’s Penguin* essay, Benkler builds a model of peer production based on institutional economics thinking. He discarded “these approaches … outside the mainstream of economic theory” in order to “establish its [peer production’s] baseline plausibility as a sustainable and valuable mode of production within the most widely used relevant analytic framework” (Benkler 2002, p. 401). Consequentially, the plausibility of peer production is based on economic calculation of a number of variables.

> “Peer production emerges, as firms do in Coase's analysis, because it can have lower information opportunity costs under certain technological and economic conditions.” (2002, p. 374)

The variables defining these conditions include: the value of property systems, the implementation costs of a property system, the latter’s opportunity costs (p. 403), costs and gains of information processing by an organisational form (407-415), “allocation efficiencies gained from the absence of property” (407, 415-423), motivations to contribute (pp.423), modularity and granularity of tasks (pp.434), costs of integration (436). In the subsequent *Sharing nicely* article on “social sharing of shareable goods” — another manifestation of peer production with slightly different underlying economics —, Benkler condensed the viability of peer production essentially to two variables. The selection of peer or open source production over market-, firm- or state-based frameworks is the result of superior effectiveness of social sharing, facilitated by favourable “motivational effects and … transaction costs” (2004b, p. 277).

The third line of reasoning refers to hindrances of peer production and possible organisational, institutional and technical means to avoid them. Discussing the viability of peer production, Benkler identifies a number of potential pitfalls and crucial elements in the production process that need to be supported either by the circumstances or by certain governance measures. The potential problems lie in the domains of provisioning, quality assurance, and integration. Economics literature has long been influenced by the meme of the "Tragedy of the Commons" (Hardin 1968) and the difficulties of overcoming the free-riding problem and attaining sustainability of commons. The "tragedy of the commons" argument states that resources held in common will eventually be overexploited and degrade as actors have incentives to appropriate common-pool resources as rapidly as possible to avoid losing out to other appropriators. Eleanor Ostrom's studies, however, provide empirical examples of how collective governance institutions can make commons sustainable and productive (Ostrom 2008; Ostrom & Gardner, 1993; Basurto & Ostrom, 2008; Feeny, Hanna, & McEvoy, 1996). Benkler describes a list of potential points of failure and their respective mollifying governance techniques or characteristics of informational goods (Benkler 2002, pp. 436-439).

For peer production to be a viable mode of production, contributors' motivations become a crucial resource. With the absence of monetary rewards and hierarchical commands, the intrinsic motivations of contributors need to be safeguarded. Based on literature of motivational economics, Benkler identified unilateral appropriation as the main motivational threat. In free and open source software projects, the General Public License (GPL), Richard Stallman's legendary copyright law hack, ensure that such unilateral appropriation is not possible, as do other flavours of copyleft licenses. A second issue that has the potential to stymie voluntary contributions is free-riding. For projects dedicated to open source software or similar informational goods, Benkler argued, that free-riding is a non-issue. While non-contributors could indeed use the peer-produced product, Benkler argued that contributors' motivation is not diminished by high numbers of free-riding users. In addition, high number of potential contributors would average out malicious inputs of dissatisfied contributors (2002, p. 441).

A third crucial element of viable peer production is a functioning integration process. The integration process assembles the actual product from diverse parts. A good integration can be achieved either by technology, iterative peer production, social norms, or even "market and hierarchical mechanisms" (2002, p. 443). Feasible ways to achieve a functioning integration comprise "technology, as with the software running Slashdot or the Clickworkers project; iterative peer production, such as the moderation and meta-moderation on Slashdot; social norms, as with Wikipedia's or Kuro5hin; and market or hierarchical mechanisms that integrate the project without appropriating the joint product, as is the case in the Linux kernel development community" (2002, p. 443). Benkler suggests that cooperative monetary appropriation would be a viable way to integrate, too. This integration problem however is only such a crucial bottle-neck problem if peer production is conceptualised as a collaborative effort with large numbers of contributors doing small tasks while only a few do some time consuming heavy-duty work. Peer production, Benkler argues, would be economically more attractive when peering of contributors is more efficient than organising their collaboration within firms or hierarchies or via market exchange (2002, p. 403). The relative effectiveness arguably increases when there is a large number of contributors only willing to contribute tiny tasks. While this argument is convincing for large-scale projects depending on large numbers of contributors such as Clickworkers and similar crowdsourcing projects, the vast majority of free and open source software projects are rather small (Schweik & English, 2007). It might certainly be more effective to have large numbers of "peer" contributors in a project, but large-group size can hardly be the decisive factor for peer production viability when the vast majority of open source projects have less than ten or even five team members.

One of the features of peer production is the self-assignment of contributors to specific tasks. In a voluntary system, any person can pick up any task at her liking. The practice of self-assignment can be seen as favourable for the motivation of contributors and for the best match between contributor's talents and the talents required for the task (Benkler 2006). However, the system fails to sort out those contributors that misinterpret their talents or have bad intentions, and therefore produce poor outcomes. Benkler argues that a high number of contributors would produce redundant results with different degrees of quality, of which the best version can then be selected for the actual product (2002, p. 380). Thus, quality is assured and the consequences of occasional defections minimised.

The fourth pillar of peer production viability consists of the norms that support the preferences of potential contributors. Norms, both in the form of regulation and cultural mentalities, can influence the behaviour of potential contributors. For example, "cultural preferences and tastes" (2004b, p. 340) are decisive when an agents has to decide whether to contribute information or work to a community of relatively unknown peers. A culture of sharing may or may not exist that would support individual contributions to a common cause. The same holds true for openness and acting transparently. Regulations and legal norms can certainly also influence the feasibility of commons-based peer production.

2.2 The Internet security governance challenge

Security and the state appear to be the least likely fields for a significant adoption of the organisational modes described in the previous sections. Traditional security institutions like armed forces, intelligence services and law enforcement units eschew public scrutiny, transparency, and public accountability. The selection of security elites, their ideational stance, their concept of security as a public good, the range of means available and the willingness to apply them is in liberal-representative democracies significantly different from what they looked like in modern state's early days under Richelieu and Fouché (Tilly 1985; Bowling & Newburn, 2006, p. 5). The history of national offices and organisations with "security" in their names is full of examples that brought anything but security to the people. Today, security institutions still are the blackest of all boxes for the democracies' electorates. A mode of production that has openness in its name — peer or open source production — hardly fits with the traditions of security politics, policies and institutions. The feasibility of the peer production of security appears low at least at first sight; capital requirements are extremely high, at least in the national security branch, and information flow and access is often strictly regulated and restricted in the security domain. Nevertheless, security governance in the last ten years has been altered by the rise of networked security production and private security provisioning.

Even more significant for the theoretical and empirical embedding of this study is the rise of significant Internet security challenges and Internet security governance over the past years. The rise of botnets, international conflicts spurred heated up by Internet-based activities, the professionalisation of Internet-based crime, scenarios, however unlikely, of the Internet as a new locus for international conflicts — all this has raised the perception that the Internet is or could become a risk, if not threat to the well-being or even integrity of actors. The rise of terms like cyberwar or cybersecurity in political discourse indicates these trends. From these problems arises a governance challenge that has been addressed by practitioners and researchers in various ways. Academics with a policy focus have contributed insights into non-technical factors that have hindered significant security improvements. Despite the lack of a range of economic incentives, the business sector has made significant investments to improve the security of its technical systems. Governments have strengthened law enforcement capabilities and international, intergovernmental cooperation; they have supported the built-up of public-private partnerships in the ICT security domain, and invested in programmes to increase the resilience of critical information infrastructures. At the same time, national security and intelligence agencies have built up cyberwarfare units, started zero-day-exploit buying programmes, and thereby decreased global Internet security. The role of social production for Internet security has not been studied thoroughly yet.

This study aims at analysing the role of peer production in the field of Internet security. The purpose of this section is to explore the domain of security, security governance, Internet security governance, and Internet security, and search the existing literature for clues as to the relevance and applications of social production or related concepts. The following subsection studies general modes of international security production and aims to locate or relate peer production to existing debates in International Relations theory, namely security governance and networks. The second subsection then dives into existing literature on Internet security, existing security problems and how practitioners and scholars have heretofore responded to them.

2.2.1 Governing security

The governance of global security is a problem that has occupied social sciences and political practitioners for the better part of the last century. The questions of the appropriate model of organising international security affects fundamental political questions and norms, effectiveness in assuring security, adequate representation and democratic rights, sufficient checks for power, and balancing legitimate political interests. And all of that on a global scale. Any model comes with its specific drawbacks. The centralisation of power — as suggested in the aforementioned global institutions — encompasses risks of potential system failure and of power abuse. There is a dilemma between the need for political institutions with global reach, and the need to avoid such institutions due to the tremendous costs of potential institutional failures (Richter 1992). Those risks alone make the idea of a world government a normatively risky institutional solution to global governance problems. The existence of a central authority does not necessarily mean that it has an authoritarian character and "attempts to be source of all rules and their enforcement" (Ostrom, Gardner, & Walker, 1994, p. 38) — but it easily could. One decisive normative task for any institutionalisation of governance of technical infrastructures is to design it in a way which prevents even inept and dishonest powers from doing much harm (Popper 1962, p. 25; Bendrath 2007, p. 12).

In theory and practice, a wide range of governance options are feasible — distinguishable by different degrees of hierarchies, coercion, shared interests, mutuality and effectiveness. Security governance is the classic topic of the studies of International Relations. Starting from the idea of international anarchy, i.e., as unregulated sphere among rival, potentially aggressive nation states, international relations theory has come up with several models to explain the absence of war. Prominent ideal-type systems for international security[17] are balance-of-power relations, col-

[17] International security is here used in its narrow sense as the absence of violent conflict in contrast to conceptualisations of international peace as the absence of structural violence.

lective security, hegemonic peace, and international regimes. This list needs to be supplemented by networked security.

Based on the construct of international anarchy, balance-of-power is the first model to provide a secure international sphere, albeit in a precarious manner. In an assumed world where individual actors, i.e., states, are not restrained by any global hierarchy or other institutionalized means, states are incentivized to maximise their influence and are even compelled to behave aggressively and therefore increase the insecurity for their peer contenders in the international arena. International anarchy forces individual states into building up their own defence, response and attack capacities. At best, the capacity build-up results in a durable balance of power, in which neither state dares to deploy force against the other for fear of retaliation from the attacked and other actors. Mutually assured destruction is the most vicious form of a stable balance-of-power formation.

Contrasting the balance-of-power model in terms of organisational precision is the collective security model, in which threats for states emerging from other states are mitigated by the establishment of a regional or global authority responsible for protecting international peace. The idea of a global hierarchy for world security was born in an intentionally mediocre way with the League of Nations or the United Nations organisation. A full-fledged hierarchical approach was designed by Grenville Clark and Louis Sohn in their book *World Peace Through World Law* (Clark & Sohn, 1958), a concept of a world government with substantial authority, a kind of super-empowered UN with a substantial executive forces to overcome the governance problem created by the invention of nuclear and hydrogen bombs. This institutional design was an attempt to overcome the "problem of discovering workable political institutions for a community … that was created by a formidable revolution in technology; … and many of its common problems are beyond the power of nation states to solve" (Barr 1958).

Hegemonic peace is comparable to collective security in that it relies on a dominant, central forceful entity, but it is a single state rather than an international organisation. The hegemon amasses power second to none, chooses to exert it only in largely benevolent ways and thereby acts as the guarantor of a hopefully just and peaceful existing order. Third states that oppose this order may face the forceful response of the hegemon, while aligned states are protected by the hegemon against attacks from third parties. The price for enjoying this gift of stable order in addition to their required support for the hegemon, however, is to endure the shortcomings of the existing order. The "benevolent hegemon" — a role frequently attributed to the United States after the end of US-Soviet conflict — ensures global security and prosperity as global public goods (Nye 1990; Mandelbaum 2006).

Using the concepts of institutional economic theory, a state's hegemony establishes a "hierarchy between polities [that] reduces transaction costs and mitigates opportunism" (Lake 2009, p. 275). A variant of such hegemonic peace is Amitai Etzioni's recently drafted concept of a "Global Safety Authority", thoroughly described in his book "From Empire to Community" (Etzioni 2004). According to his observations, a supranational police agency is on the rise: "what was going to be the new global order following the collapse of communism and the bipolar world, has been answered following 9/11… It has basically taken the form of a global police department that I refer to as the Global Safety Authority (GSA), which is run by the United States and its allies but encompasses most nations of the world. … As a result, the intelligence and police services of scores of nations work together quite seamlessly, sharing information, covering actions and even integration" (Etzioni 2005, pp. 472-473). Etzioni's observations serve as a reminder of the more than a decade-old debate among US foreign policy pundits about an "information umbrella" as the strategic means and brace for perpetuating the US-led post-World-War II alliance (Nye & Owens, 1996; Schmidt 2004). The rise of a network of ever more interlinked national Computer Emergency Response Teams (CERTs) and especially the transnational cooperation among national intelligence agencies could be seen as a de-facto global Internet security authority in its early stages, as early steps towards an "information umbrella."

The fourth fundamental way to ensure a peaceful international order mixes some of the characteristics of the previously described approaches. Lacking stable orders provided by the models of collective security or hegemonic peace, states can still reduce their mutual distrust by establishing international regimes and norms. States can manage to balance their security interests, reduce mutual distrust and establish an international order that does not resemble a zero-sum game. For many international issues, international regimes are the default organisational form of problem solving.

Resembling international regimes, networks in various forms have entered the sphere of global politics as an organisational form. The concept of "transgovernmental networks" (TGNs) reflects the widening and deepening of international collaboration, and the intensification of communication between the medium and lower levels of national bureaucracies (Raustiala 2002; Slaughter 2004). These TGNs manage to produce outcomes beneficial to the states involved. During the last decade, security and policing studies have observed a diversification of how security is provided, away from the state as the sole provider of public security towards a system in which the state is supplemented by private actors such as security services and mercenaries. In national security circles, the term "networked security" refers to "loose institutional arrangements and non-hierarchical structures of information exchange" (Gruszczak 2008) that are established, for example, in anti-terrorism activities or to re-establish security in former failed state such as Afghanistan (Jung 2009). However, the idea of networked governance goes beyond the idea of networks as a governmental tool.

In international relations, the state has long been portrayed as the dominant political player. Increasing density of international institutions and intergovernmental cooperation has led to rise of scientific models, theories and concepts such as institutionalism, global governance to provide scientific tools for analysing and assessing recent institutional and organisational developments in global policy fields. But it appears as if the institutional landscape has changed even further, and led to the rise of the term and concept of networks to explain frame social realities: "networked politics" (Kahler 2009b, 2009a), "government networks" (Eilstrup-Sangiovanni 2007), "transnational networks" (Eilstrup-Sangiovanni 2005), "network approach" (Dunn-Cavelty & Suter, 2009), "network governance" (Sørensen & Torfing, 2007), "networked governance" (Vymětal 2007; Parker 2007), "nodal governance" (Shearing & Wood, 2003), "multi-nodal politics" (Cerny 2007), "networked polity" (Ansell 2000), "network of terror" (Mayntz 2004) — all of that in the "networked century" (Slaughter 2009) of "global government networks" (2004). As Miles Kahler puts it: "Networks have become the intellectual centerpiece for a new era" (2009b)—that of "networks and states" (Mueller 2010). The question remains however: which kind of networks?

The value of network analysis for students of international politics lies in its ability to provide new perspectives on another dimension of power, based on relationships and the position of an actor within a network, and that it frames the congregation of networked actors as a distinct actor on its own. Networks thus allow for new power structures and new methods of empowerment, that can both be used by these new "networks-as-an-actor" (Kahler 2009b) or by conventional actors such as states. A conventional neorealist theory of global politics hardly allows for a proper representation of networks-as-actors and can create those pitfalls of oversimplifying and overestimating the role of governments.[18] For Anne-Marie Slaugther (2004), it is a network of jurists, parliamentarians, regulators and other experts that link up with their peers in other countries, forming a "new world order" with "disaggregating states". With governments as "aggregations of distinct institutions with separate roles and capacities" (2004, p. 13) and deconstructing the state into its functional pieces, the new "Lego world" (Slaughter 2011) brings forward a new form of governance that is characterised by "regular and purposive relations among

[18] Cf. the critique on Goldsmith and Wu (2006) by Cowhey and Mueller (2009), FN 20.

like government units working across the borders that divide countries from one another and that demarcate the 'domestic' from the 'international' sphere" (2004, p. 14). While Slaughter acknowledges transnational networks as actors in international politics, her analysis falls short of providing a framework for analysing those networks providing governance functions that have emerged out of distributed collaboration.

Global governance discourses have ever centred on the role of states in future governance structures. Many studies support industry-based self-governance and other forms of regulation by non-state actors (Hutter & others, 2006; Dilling, Herberg, & Winter, 2008), while others have shown hints of inefficiencies in privatised operations of once public infrastructures (Groenewegen 2005). It is not just a question of efficiency though, but also one of general capability. Furthermore, societies are confronted with problems, the causes of which and, more importantly, mitigating means are beyond the territories and occasionally also the capacities of single states. Global warming is one of those issues, the Internet and problems related to it another.

Both in security studies — based in International Relations and dealing with questions about international security, peace and war — and in policing studies, there has been a trend in recent years, to decouple the provisioning of security from the existence and actions of the state (Bayley & Shearing, 2001; Bowling & Newburn, 2006; Bryden & Caparini, 2007; Caparini 2007; Dupont, Grabosky, & Shearing, 2003; Dupont 2004; Hänggi 2003; Huysmans 2006; Kempa, Carrier, Wood, & Shearing, 1999; Krahmann 2003; Zedner 2003; Felício 2007; Waugh, Jr., & Sylves, 2002).

For researchers in the field of international relations, security governance is not just a term that labels the governance of the security sector. The new realities of the post-1989 world called for a new term to describe "this delegation of authority and the outsourcing of public policy functions" (Krahmann 2003, p. 11). But security governance is also used to label a framework to analyse security politics in the post-Cold War world. The core features of this framework are its inclusion of non-state actors, that it considers new governance structures such as networked cooperation, uses a broadened and widened concept of security and thus of security areas, and concentrates on security provisioning processes instead of focusing on security organisations and players (Krahmann 2005, 2003). Further elements of this overall 'pluralisation of security' processes are the "increased emphasis upon 'high policing'", "changed roles of law enforcement and security agencies", "a blurring of the boundaries between international security and domestic concerns of order maintenance", and a new idea of coercive policing accompanied by "securitization and militarization" of police forces and increasing convergence of police, military and intelligence (Bowling & Newburn, 2006).[19]

2.2.2 Internet security

Discussions of Internet security and the appropriate institutions to contain existing or emerging risks and threats are filling the conference rooms and meeting tables of academics, policy makers, and cyber intellectuals. The debate has been spurred by countless minor security incidents resulting in data leakages and system intrusions and some landmark high-level incidents such as the Estonian cyberattacks, the Conficker botnet, Stuxnet, and diverse attacks on U.S. ICT systems attributed to Chinese actors. Internet security has become a hot political issue on global scale, leaving behind a past in which it was only of interest to geeks, technical regulators, or criminologists. Policy makers and national bureaucracies have reacted to rising numbers of Internet security incidents and the perception of increased vulnerabilities. Political response to these Internet-based threats, risks and vulnerabilities combined with the diagnosis of insufficient means against them has been a mixture of increasing public awareness, fostering private self-regulation and public-private cooperation, creating Internet security groups within traditional state-based security organisations, and fostering international incident-response exercises. Over the past years, traditional security institutions such as military, intelligence and law enforcement have apparently increased their attention on questions of Internet security (Anderson 2013). Every other week, another country updates its cybersecurity strategy, proposes new Internet security-related legislation, or sets up yet another cybersecurity initiative ("UK Launches New Cyber Security Strategy," 2011). The range of security risks is obviously as wide as the range of possible responses to them. This section therefore seeks to portray the scientific and policy landscape responding to the Internet security challenge.

First though a few details on the conceptualisation of Internet security used throughout this study (Schmidt 2009). In many studies, 'Internet security' is often referred to as the absence or the proper handling of 'security problems' like phishing and spam, and achieving the technical resilience of the Internet (Anderson & Moore, 2007; Brown & Marsden, 2007). Threats to Internet security are often described by terms like 'cyberattack', 'cybercrime', 'cyberterrorism', or 'cyberwar'

[19] Cf. also definitions of governance by Caparini 2007, p. 269, and Krahmann 2003, p. 11. Noteworthy: "…governance can be differentiated from government along seven dimensions: (1) geographical scope, (2) functional scope, (3) distribution of resources, (4) interests, (5) norms, (6) decision-making and (7) policy implementation.“ (2003, p. 12)

(Bendrath 2003; Dunn Cavelty 2007; Wilson 2007). Others prefer 'cybersecurity' over 'Internet security' (International Telecommunications Union 2005). While some link cyberwar to state actors by defining it as "deliberate disruption or corruption by one state of a system of interest to another state" (Libicki 2009), others prefer a broader concept of cyberwar, which includes any type of actor. These examples indicate that the term 'Internet security' suffers from the same problem that Myriam Dunn Cavelty has identified with regard to the term 'cyberterror' — it's "a very elusive and poorly defined concept" (Dunn Cavelty 2007, p. 22); and these definitions do not meet the elaborated criteria set for conceptual explication in general.

So far, conceptualizations of Internet security or similar terms have failed to draw on the theoretical insights of other social scientific streams analysing 'security'. A few of these various facets shall be listed here. First, according to constructivists and Copenhagen School aficionados it is impossible to intersubjectively define security, which leads to the so-called securitization model in which "security is what a political actor or a political entity labels as security in a particular situation" (Daase 1993, p. 45). Second, security is a recursive concept; security is "the assuredness of the reliability of protection or risklessness and the hence resulting state of unconcern" (Kaufmann 1973, p. 344). Hence, it describes not only a security object, but also the means to protect that endangered good. Third, Internet security not only has political and psychological dimensions, but a technical one, too. 'IT security', which these days is basically synonymous with Internet security,[20] is comprised of three fundamental principles: availability, confidentiality and integrity of data (Eckert 2001). Fourth, states of technical security are not necessarily congruent with an actor's security priorities. Technically insecure ICT systems, i.e., low levels of availability, confidentiality and integrity of data, do not always pose a general security risk for users; it depends on the usage and the social, political and economic functions of the technology. Last but not least, Baldwin identifies six dimensions or "specifications" of security "that would facilitate in analysing the

[20] This claim was made by Christoph Meinel, professor for computer sciences and CEO of the Hasso-Plattner-Institute, Potsdam, Germany.
Cf. http://www.hpi.uni-potsdam.de/meinel/teaching/lecturesclasses/Internetsecurity_bjut.html for his
presentations and online lectures, last accessed in May 2010.

rationality of security policy": security beneficiaries, security objects, degree of security, security threats, security means, costs of security and security timeframe.[21]

To conclude, a conceptualisation of Internet security that factors in the aforementioned findings is necessarily rather broad. It would result in the following conceptual beast: Internet security is the low probability of damage to acquired values forming the Internet (such as sub-networks, computing devices, components; integrity, availability, reliability of data) or based and depending on the Internet (such as contents and semantics; economic welfare, autonomy, political independence), which are related to beneficiaries (such as individuals, states, social groups, businesses), with the aforementioned low probability of damages achieved by applying protecting means (either technical, organizational or political in nature) against threats (emerging from either malevolent or ignorant actors, from systemic constellations, technical phenomena or artefacts). In shorter and plainer language: Internet security is when the things an actor values are likely to be fine now and in the future and not be harmed by anything related to the Internet. The empirical analysis in this study focuses on a subset of all possible instances of Internet security cases, namely on large-scale incidents that endanger the very technical functionality of parts of the Internet's infrastructure.

The design of Internet security governance and production depend on countless factors. Their complexities and interconnectedness turn the institutional design of Internet security into a wicked game for a number of reasons. Firstly, security policy usually involves force, enforcement, and some degree of secrecy. Second, the design of working Internet security institutions is a transnational task. The distributiveness of security problems, of incidents, of systems involved, of perpetrators or attackers, of actors required for mitigation, require global solutions and a distributed organisation. Third, Internet security mingles foreign with domestic security, and foreign policy with public policy. The practices of foreign and national security have traditionally differed from those in the domain of homeland security. Applying the former to the latter substantially alters the latter and long established societal norms. Fourth, all these factors combined have a great potential to make precarious the legitimacy of Internet security policies.

Internet security governance and production has been and still is based on a mixture of different organisational approaches. They range from the peer production

[21] These seven nouns paraphrase the seven questions Baldwin has raised to identify the seven dimensions of security: 1) Security for whom? 2) Security for which values? 3) How much security? 4) From what threats 5) By what means? 6) At what costs? 7) In what time period? (Baldwin 1997, pp. 12-17)

of security standards at the IETF, to market-based exchange of security products and services, to forms of industry self-governance, to networked organisations exemplified by the anti-spam London Action Plan (Tabatabaie, van Eeten, & Asghari, 2012), to various public-private partnerships on regional, national or international levels, to international treaties like the Council of Europe's Convention on Cybercrime (2001), to initiatives of international organisations like the ITU, Interpol's Global Complex for Innovation, or NATO (Tikk 2010), to intergovernmental cooperation among law enforcement, intelligence, and military units, to unilateral approaches. Comparing the Internet security governance status quo with the models of global security portrayed in the previous section, the status quo apparently does not fit to just one ideal-type. On the international level, there is neither an established and globally accepted cyber hegemon, nor do collective security institutions exist (Schmidt 2014). With regard to policing, national law enforcement units have gradually increased their attention and competencies (Anderson 2013), while national laws are harmonised by international treaties.

A unifying feature of all approaches however has been the need for private-public cooperation. States, the actor whose core feature it is to provide public security, typically do not own and have little direct control over the technical systems that make up the Internet. Private companies, which usually do own these components, have economic interests that are not necessarily congruent with the public's need for security. States have responded to these challenges by setting up public-private partnerships (PPP) as the perceived "panacea for this problem" (Dunn-Cavelty & Suter, 2009, p. 179). Both in Europe and the US it appears to be the preferred governance mechanism. The EU has followed the idea of a European PPP for resilience (EP3R), a project started in 2009 (A thorough analysis of EP3R: Irion 2012). A new European directive on network and information security that is currently debated in EU institutions would mandate the sharing of certain information and grant as yet unestablished national "competent authorities" some authoritative control over attacked ICT systems (European Commission 2013). The resulting security architecture would hence have more public elements. Dating back to the days of the Clinton administration, an increasing number of private-public partnerships on all governmental levels have been set up in the US to enable a seamless flow of the information required to respond to incidents effectively (Nyswander Thomas 2013, pp. 9-14). Transferring the public-private partnership approach to Internet security is problematic, though, as scholars of organisational aspects of Internet security such as Dan Assaf (2007), Amitai Aviram (2004), Kristina Irion (2012), Myriam Dunn-Cavelty and Manuel Suter (2009) attest to. The inherent clash of underlying goals of PPP with the goals of any cybersecurity arrangement is an obvious example. The former strive for and promise efficiency and involve partners from the private sector, the ultimate economic interests of which usually do not result in the production of a public good. This convergence problem has raised doubts as to whether PPP is the best organisational approach to Internet security. (For more details on this debate cf. Schmidt 2013b)

A great chunk of the non-computer science research on Internet security has focused on the economics of information security. In his article "Why Information Security is Hard — An Economic Perspective", Ross Anderson argued that "information insecurity" could be "explained more clearly and convincingly using the language of microeconomics: network externalities, asymmetric information, moral hazard, adverse selection, liability dumping and the tragedy of the commons" (Anderson 2001). This increasingly rich body of literature has provided analyses on the incentives of various actors to intervene or ignore ongoing security problems — actors that are affected by security problems, own or operate Internet connected systems, produce hardware or software components, provide relevant services, provide policing services or have the ability to alter the regulatory environment (Anderson & Moore, 2006, 2007; Anderson, Böhme, Clayton, & Moore, 2009, 2008; Tsiakis & Sthephanides, 2005; Ioannidis, Pym, & Williams, 2009; Bauer, van Eeten, & Chattopadhyay, 2008; Bauer & van Eeten, 2009; Moore & Clayton, 2008; Moore 2008; Moore, Clayton, & Anderson, 2009; van Eeten 2008; van Eeten & Bauer, 2009, 2008b, 2008a; van Eeten, Bauer, Asghari, Tabatabaie, & Rand, 2010; van Eeten, Asghari, Bauer, & Tabatabaie, 2011; van Eeten, Bauer, Asgharia, Tabatabaie, & Rand, 2012).

These studies have shown in great depth how misaligned incentives for owners of certain types of ICT infrastructure have helped to increase insecurities or at least stymied possible interventions to decrease them. Some actors own and operate Internet components, but are hardly economically affected by insecurities and therefore have very little incentive to intervene.

International relations scholars have added a significantly different perspective on the Internet security problem with their traditional focus on international order, relative losses and gains of power, hegemonic stability, or deterrence, to name a few core concepts. From this perspective, Internet security is not an issue of global public policy, but an indication that the Internet has become a sphere contested among nations, and a resource of national power. Drezner argues that global governance would primarily be driven by great powers and multipolar regulatory divergence: "even when states prefer to let private actors take the governance lead, they will intervene at crucial junctures to advance their desired ends" (Drezner 2007, p. 118). Internet politics would eventually resemble conventional geopolitics (Drezner 2007, 2004). Similarly, Goldsmith and Wu assessed that the Internet would become an arena for traditional great power rivalries (Goldsmith & Wu, 2006, p. 184), in which national governments would establish themselves as the driving actor in Internet governance. Mueller however, argued that there would be "no simple reassertion of the national state and reigning in of the global Internet" and highlighted the "virtues of the softer, networked methods" of Internet security provisioning (Mueller 2010, p. 183). Nye assessed the applicability of deterrence as a governance technique to ensure Internet security, and concluded that "some of the most important security responses must be national and unilateral" (Nye 2011a, p. 36).

This section has outlined a wide range of aspects of Internet security and possible approaches to respond to a diverse set of security challenges. The study touches on several aspects of the existing research on Internet security described above. Its theoretical foundation, however, is theories of open source and peer production. Literature on Internet security and Internet governance has mentioned the operational Internet security community at various instances. Nevertheless, systematic, empirical research on transnational Internet security incident response has effectively been absent at the time this study commenced. Though there are first publications on the history of cyberattacks, a history of responses to cyberattacks is still due.

2.3 Secrecy vs. openness

A main theoretical problem that drives this study is the apparent clash between secrecy and openness and its consequences for the viability of peer production in the security domain. The clash is not constrained to the mere conceptual level, but resides deep within different theoretical stances in various academic disciplines. For cryptographers, methods and implementations need to be open. Any technology, algorithm and software that is involved needs to be transparent to ensure the security of, say, encryption keys. Operational data such as pass-phrases however must be kept secret. For national security planners, any strategy, measure, tool, and resource involved in securing an "acquired value" like, say, a country, against foreign aggressors is best kept in the dark for anyone but a handful of people. For police officers involved in investigations against organised crime or others as aggressive as capable adversaries, operational security is a prerogative to protect their work and their personal well-being. For those German civil servants investigating tax fraud within highly influential organisations that were later declared mentally disordered and removed from their posts (Hunfeld 2008), openness towards the public might have served as a protection against political scheming.

The literature on secrecy comprises such different topics as secrecy in national security investigations (Sales 2007), the role of disclosure for security provisioning (Swire 2004), and general discussions on government secrecy (Roberts 2006). The role of secrecy, its necessity, consequences and price for Internet security governance appears not yet to have been covered.

As described, the reason for and against secrecy and openness are manifold. This section seeks to portray these perspectives and relate them to the question, which impact secrecy could in theory have on the viability of peer production. The argument is that some degree of secrecy is likely compatible with the peer production model, but that the application of secrecy alters the underlying economics of the peer production model. Secrecy probably a) makes the model less viable, b) potentially undermines the ideational societal effects of peer production, and c) likely results in a different kind of social production.

Before going into the details in the following subsection, the core concepts of this section need to be clarified: secrecy and openness. Secrecy could be defined as denial of certain information or access to certain information to a set of actors. In the context of products and their production, openness of a product refers to its accessibility and modifiability (cf. following paragraph), i.e., whether an actor is granted the right and possibility to a) use the product, b) access the source code, the recipe and the ingredients for the production process, or c) alter the product by altering the recipe.[22] Consequentially, secrecy is the counterpart of openness. Secrecy describes the intended refusal of actors to grant another actor access to information. Openness refers to the accessibility of information. Closure is accordingly used in a synonymous way; disclosure describes opening information, i.e., making it accessible.[23] Another aspect of openness and secrecy is noteworthy. While at first sight the opposite appears to be true, openness and secrecy are no binary concepts. “They need to be placed on a continuum that ranges from closed to open and encompasses varying degrees of openness.” (2006, p. 122) Referring to systems, such con-

[22] Maxwell differentiates between two dimensions of openness. Availability and accessibility describe whether information and the results of knowledge production are available for others. Responsiveness and modifiability refer to whether the results of knowledge production are modifiable by others (Maxwell 2006, pp. 122-123).

[23] One could argue that disclosure is the more appropriate counterpart to secrecy, as disclosure would include an active intended activity to grant access to information. However, openness does not depend on disclosure inasmuch as the information could have been undisclosed from the start. More so, secrecy is both an act of keeping information secret as it describes a situation in which information is kept secret. Given this close semantic relationship between secrecy and openness, it comes not as surprise that discussions on either have some similarities.

tinuums certainly exist. With regard to one particular set of information, openness and secrecy however are binary concepts. Openness of a production system therefore is more a continuous, rather than a binary concept. [24]

2.3.1 Functions of secrecy

In social, political and economic relations, secrecy can have various functions. On the political level, general secrecy helps to protect the interests of an actor. It is the kind of interests, the interested actor and the excluded actors that form the legitimacy of secrecy from a democratic governance perspective. For an actor, the application of secrecy is a utilitarian decision, resulting from the advantages and disadvantages of applying secrecy or openness.[25] Therefore, secrecy can have a different function. It can be in the public's interest, a means for bureaucratic cover-up, a principal's instrument, an indication of asymmetric power relations, an underminer of legitimacy and a democratic public sphere, and instrument to gain competitive advantages in business matters or international affairs.

Security in modern nation states has traditionally been in the hands of dedicated authorities, of which a distinguishing feature is the ubiquitous application of secrecy and non-free flow of information. The legitimacy of secrecy as it is applied in traditional security organisations is fed by the nature of the excluded actor — criminals, spies, foreign armed forces —, the risks inherit it disclosing information to them, and the idea that these organisations serve a common interest and provide security as a common good. In settings in which insecurity is caused by actor-based threats, secrecy is a technique to deny adversarial actors a tactical advantage based on an informational edge. When a society is beleaguered by enemies prying for any kind of security-sensitive information, secrecy can at times be defined as a "collective good". As William Colby puts it: "Secrets are necessary to a free society" (Colby 1976, p. 4), irrespective of the fact an autocrat like Stalin relied on secrecy for his very survival, too (Harrison 2003, p. 3).

However, the realities of public authorities and bureaucracies recommend a different perspective on secrecy. The issue with public bureaucracies is that they develop a kind of *eigenleben*, an independent existence. According to Max Weber, secrecy is

[24] This categorisation is still rather coarse, only includes results of knowledge production and lacks to incorporate processes and the design and governance thereof. In its current state, the scientific discourse does not provide insights into more finely grained aspects (processes, knowledge objects) of Internet security governance and their relationship with openness respective secrecy.

[25] More on that in the following subsection.

a way for bureaucracies to amass power, to evade public scrutiny, to avoid accountability and to conceal controversial activities (Weber 1921/1972). Secrecy is a technique not only to conceal power relations, but also to implement them.

According to Mark Harrison, secrecy is a means to ensure the implementation of commands in hierarchical systems. In hierarchies, commanding principals face the problem that agents might not obey unconditionally, but only "conditionally, or shirk, or steal the principal's advance and invest it in an exchange with an external private network" (Harrison 2003, pp. 9-10). A principal has several techniques on hand to solve this command-obedience problem. Historically, the problem of command has been solved by a variety of mechanisms including monitoring, "internalisation and promotion, rewards for compliance, the penalisation of shirking and disloyalty", the use of force to minimise rewards for a disobedient agent and his collaborative network (2003, p. 11). Secrecy is another governance technique a principal can apply to drive up costs for deviating agents.[26] The implication of these ideas is that openness is not per se an indication of a benevolent principal. It could just as well mean that she has different means to deal with his command-problem, does not have the means to revert to secrecy, or the application of secrecy has effects detrimental to the principal's interests.

Nevertheless, there are strong normative arguments against secrecy as a utilitarian instrument for security. Liberal societies and democracies depend on openness for the electorate to judge the performance of the elected (Gowder 2005), and therefore secrecy inevitably undermines democratic principles. This effect nurtures an ethical position that assumes "that a wrong is a wrong, regardless of the greater harm it avoids" (2005, p. 25). One of the requirements for the possibilities of a civil public sphere, to use Jürgen Habermas' terms, is the absence of *Arkanpolitik*, arcanum politics (Habermas 1962/1990, p. 120), which "institutionalizes the *de facto* unreviewable security choices of powerful elites" (Gowder 2005, p. 1). Viewed from a normative democratic perspective, secrecy will ultimately lead to regulatory and governmental failure.

The observation that secrecy undermines the public sphere is stressed by Edward Lee. The political effects of secrecy become obvious when looking at the political effects of openness, as they materialise in the concept of the public domain. In his "unified theory of the public domain", Lee conceptualises the public domain as a

[26] Harrison assumes that under secrecy the exchange of certain information is illegal; hence any private transaction using this information becomes illegal; with the illegality, an actor loses the ability to enforce failed transactions (2003, p. 13).

means to impose “legal restraint on the government’s ability to restrict the free flow of information, ideas, or materials” (Lee 2003, p. 97). The legal mechanism, applied in areas such as intellectual property law or government secrecy, is to "[accord] each person a right of unrestricted access to or common ownership of material in the public domain” (2003, p. 98). By imposing secrecy on certain matters and information, a government can grant access to resources to a limited number of actors (or just one actor: a monopolist). With the legal institution of the public domain, “an ultimate restraint on the government's power” (2003, p. 118) is established. In short: secrecy is power, openness restricts power. To sum up, there are several readings of secrecy and its relation to technical and functional security, its normative and political effects. A strong political trend however points into the opposite direction. “The moment we now face is, by all measures, democratic-disabling: The government asserts the power to shrink the public domain (by removing thousands of works from the public domain) and the power to skirt it (by shifting its conduct from the purview of the public domain to the domain of secrecy)" (2003, p. 209). In conclusion, secrecy involves a number of normative and organisational consequences that should make decision makers think twice about its application. The problem of secrecy becomes obvious in the light of the characteristics of openness.

Secrecy is not only a means to achieve political advantages. Probably as important is the role of secrecy in everyday business conduct. The at times legal, at times social institutions of proprietary information[27] plus a variety of technical and organisational precautions help to secure the “secret sauce” (cf. the quote cited by Swire in this section) of market organisations and ensure an organisation’s competitive advantage. For any actor contemplating the disclosure of certain information, there are a variety of variables to take into account (cf. the subsequent section). In a rational-choice calculation, the decision comes down to two decisive factors: “[T]he incentives of disclosure depend on two largely independent assessments — the degree to which disclosure helps or hurts security, and the degree to which disclosure creates competitive advantages or disadvantages for the organization” (Swire 2005, p. 1335).

[27] Discussing the “Disclosure of nonproprietary information” in the business sphere, Ronald Dye defines “proprietary information … as any information whose disclosure potentially alters a firm's future earnings gross of senior management's compensation. … This includes information whose disclosure could generate regulatory action, create potential legal liabilities, reduce consumer demand for its products, induce labor unions or other suppliers to renegotiate contracts, or cause revisions in the firm's credit standing in addition to that information which is, in the traditional sense, strategically valuable” (Dye 1985, p. 123).

Practically applied, however, secrecy can at times turn out to be a hindrance to efficiency. A telling anecdote is the secrecy hiccup during the 1947 secrecy reforms in Stalin's Russia. Subordinate staff could not be informed about their new responsibilities because the information was not allowed to reach them according to the old secrecy regime (Harrison 2003). In his brief article on the sociology of secrecy, Ritchie Lowry discovered three important functions of secrecy that are beyond the aforementioned rationalisation of secrecy: a) secrecy assures the "production and protection of relatively useless and unreliable knowledge", b) it provides "the consequent guarantee of individual and organizational job security, c) it results in an "extension of secrecy into areas involving sensitivity" (Lowry 1972). New innovation and production models, however, depend on free flow of information and unclassified data.

The function of secrecy as a source of competitive advantages is well known in politics, too. In international politics, gaining relative competitive advantages and hindering other states from achieving them are the driver of national foreign policies, at least in neorealist and in institutionalist school thinking. While Swire differentiates between security interests and competitive advantages, these concepts become more blurry in the domain of international politics, where the factual definition of what constitutes a national security interest and a competitive advantage tend to overlap.

The idea of secrecy as competitive advantage is also well known in circles connected to open source software. Secrecy and open source are not mutually exclusive, first of all among users, but also among the producers. The openness of open source software is usually protected by technical and legal means such as the GPL, which require the disclosure of the source code with the distribution of the software. Nevertheless, even the relatively strict GPL still allows actors to secretively use, adopt and change open source software as long as they do not distribute it (Swire 2005, p. 1353). One of the key incentives for secrecy for both individual and organisational developers of open source systems is to "stay ahead of the curve", i.e., their competition (2005, p. 1356). Corporations involved in open source software have strong incentives to keep their recipes hidden regarding how they use, integrate and adopt open source software to develop new services and business models. The virality of GPL software only infects the source code of distributed software with openness, not entire organisations which consume or fiddle with it. It does not spill over to configuration files, integration techniques, and business models.

Swire cites Robert Lefkowitz, former director of Open Source Strategy at Merrill Lynch: "You can take that open source stuff, add your secret sauce, and create some very nice commercial products." (2005, p. 1357) This usage of open source software might create conflicts between users and creators of open source products, as it might be perceived as an illegitimate appropriation of common good. The GPL and secrecy here help to secure business models and to secure the crown jewels of a number of enterprises with information systems based on open source software. Google, the company usually evangelising open technologies — "Open usually wins", as Google's then Android boss Andy Rubin put it (Stone 2010) — is closed when it comes to their cash-cow, the advertising business. "In other words, companies are very closed, secretive, and controlling about the part of their business that makes the money." (Elgan 2010)

But secrecy is woven into the open source software ecosphere not only for commercial purposes, but arguable also for security reasons. Swire states, "secrecy…is used to help security for Open Source software." (Swire 2005, p. 1347) The examples of secrecy mentioned by him include keeping password and encryption keys secret,[28] the surveillance of attackers, and the shielding of information regarding the use and configuration of defence systems. The latter include intrusion detections systems, honeypots, and firewalls. They are used, in this context, to protect open source production environments. Key rationale for defenders is that the disclosure of such information would help attackers more than it would support the defensive side. Swire concludes: "Within the overall paradigm of Open Source software, which so emphatically emphasizes the security advantages of disclosure, there are situations where secrecy appears to improve security." (2005, p. 1352)

2.3.2 Secrecy and technical security

Of particular interest in the context of this thesis is the relationship between secrecy and security, especially the effect of secrecy and openness on technical security. Among cryptographers, the question as to whether openness of source code and algorithms favours the attacker or the defender has been hotly debated since the nineteenth century (Anderson 2002, p. 1). In computer science and security economics, the question is still looming as to whether openness fosters security or obstructs it. The answers of scientists on this, in short, are usually akin to "it depends". Fundamentally, there are two perspectives, nicely summed up by the two claims cited above: "no security through obscurity" vs. "loose lips sink ships". The former represents the common thinking among open source software developers;

[28] Not taking conventional wisdom and practice for granted, Swire indeed discusses the possible gains of disclosing passwords and sharing secret keys.

the latter view is shared by the military and intelligence communities. Sociologists and political scientists have discussed secrecy, openness and their relationship with power, democracy, freedom and security politics.

The relationship between secrecy and openness on the one hand and security on the other has been addressed by a number of researchers. Computer scientist and security economist Ross Anderson's answer to this is simple: “In a perfect world, and for systems large and complex enough for statistical methods to apply, the attack and the defence are helped equally.” (2002, p. 10) However, this symmetry can — according to Anderson — be broken by a number of information asymmetries, principal-agent problems and incentive structures, such as transaction costs that favour a vendor to remain closed, regulations mandating exploit disclosure to national intelligence agencies, time-to-market for patch releases after exploit disclosure and different incentives among a group of defenders occasionally causing free-ride behaviour. Hoepman and Jacobs assent to this judgement. Openness could increase the exposure of a system. Exposure, defined as the likelihood of a successful attack, depends on factors “like the number and severity of vulnerabilities in the system, … whether these vulnerabilities are known to attackers, how hard it is to exploit such a vulnerability, and whether the system is a high-profile target or not” (Hoepman & Jacobs, 2007). Hence, regarding technical security, the answer is a strong ‘it depends’. These two different stances can be observed in any of the frequent debates about whether particular security-related information like vulnerabilities, incidents or ongoing attacks should be disclosed or not.

As to the question above, whether “no security through obscurity” or “loose lips sink ships” hold true, whether open or closed software is more secure, Swire takes a position in between. “Disclosure is often additionally useful for promoting long-run security and assuring accountability.” (Swire 2004, p. 33) With his *Theory of Disclosure*, Swire provides the tools to analyse costs and benefits of disclosing or hiding information affecting security. Whether openness fosters security in the long run depends on a number of variables and contexts that vary from case to case and make it thus impossible to assume that hiding one's security relevant information deteriorates security.

Among these *variables* are the organisational locus of relevant expertise (inside/outside), the incentive structures to improve the defence; the persistence of relevant expertise (vendors can disappear from the market, but organisations have budgets to fund positions); the institutional context for patching and fixing security issues; the effectiveness of institutional mitigations (audits, inspectors, oversight boards) against failure-proneness of secretive environments; (2004, pp. 30-35); the effectiveness of a first attack (high risk of success supports hiddenness); the number of attacks (high numbers of repetitive attacks favour openness); the ability of learning from failures by attackers (high ability of learning favour closure); the degree of communication and information sharing of lessons learned among attackers; and, last but not least, defenders' ability to alter their defence (2004, p. 12). Cryptography-expert Whitfield Diffie adds another variable: clarity of separation between always-hidden key and the usually public parts of a system (2005, p. 1348).

A good summary of what utilitarian thinking about secrecy for security is provided by Paul Thompson. In the context of Internet and security, Thompson identifies three reasons for actors to keep information secret: a) to not provide maliciously intended attackers with ideas that might support them in their activities, b) to protect information for personal security, c) to ensure technical security to ensure Internet functionality, personal security and public security (Thompson 2000).

2.3.3 Peer production, openness and secrecy

With the functions of secrecy and its effect on technical security clarified, the spots are now again on the other theoretical main topic of this study, peer production, and how it relates to secrecy. The assumption is that security production requires some degree of secrecy, that peer production requires openness, and therefore the two of them do not get along very well together. Peer production, as we recall, is based on the definitional characteristics of distributiveness, socialness, and openness. This subsection describes how secrecy and the limitation of openness are considerable obstacles for the viability of peer production in several respects.

Peer production consists of three dimensions (cp. Figure 2.2): first, the input resources that are required in the production process; second, the production platform, i.e., the organisational and technological base where contributors gather and collaborate; third, the product dimension which describes intermediary and final outcomes of the production process. In the production dimension, secrecy/non-openness affects the accessibility of information necessary for production, the number of contributors, motivation of contributors, overall productions costs and so on. In the product dimension, non-openness affects the access to the product, and the right to alter and innovate based on the product. Input resources are affected by non-openness in similar ways.

Before going into detail, a comment on the conceptualisation of openness. As the discussion in other sections has shown, the openness of a production platform is not as binary a concept as it might seem at first sight.[29] To label something as open or not in the social world, is a decision with an arbitrary element. There might be reasons to call software like, say, Android an open platform just like there are reasons to call it a hybrid of openness and closure. While the Android code base is accessible and available as a downloadable SDK, the main code base is governed by a single company and protected by a variety of contractual and organisational control points (Constantinou 2010). The decision whether to call hybrids open or not needs at least to be apparent and justified. However, the decision to call a hybrid platform that is characterised by a mix of open and closed elements an open platform leads to the interesting side effect that secrecy and openness appear to be compatible. This would apparently lead to conceptual mess and therefore this study tries to use a stricter definition of openness.

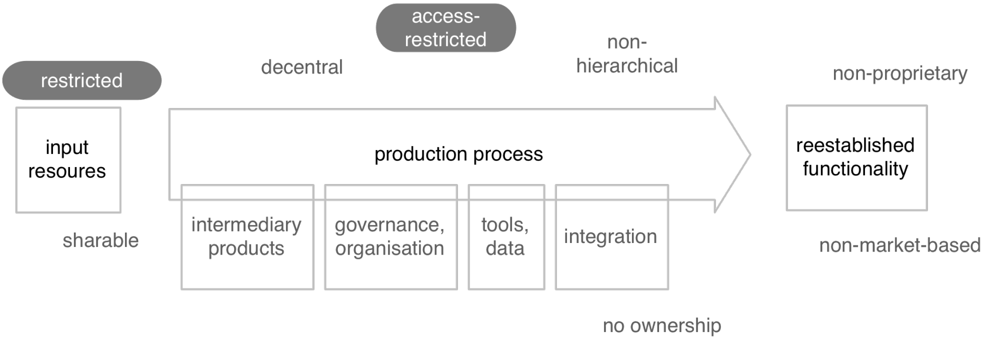

Figure 2.2: Impact of secrecy on availability of resources and (intermediary) results, and accessibility of the production platform

According to its defining characteristics, peer production requires non-restriction or openness in the sense of access to and modifiability of production resources, production platform and products. Its twin concept *open* source production even has openness in its name. The theoretical impact of secrecy on peer production can be categorized along these three elements of the production process: input, the production platform and the results. First, the impact of secrecy on the input resources. Information production requires other information and knowledge as a production resource. As peer production is by definition not based on market mechanisms, these resources need come without a price tag. In its core model, peer production requires a potentially large, undefined number of persons contributing to knowledge production. Contributors need free and unhampered access to information and intermediary products. Limitations to access or availability of input

[29] Cf. sections 2.3, Secrecy vs. openness, and 2.1, Social production and the collaborative turn.

resources could come in the form of legal, contractual or social provisions. Restrictions of availability or freedom of use of input resources would serve as monetary or motivational deterrent for potential contributors as they could increase transaction costs for a potential contributor. Furthermore, low availability or accessibility of high quality information might decrease the quality and value of the product resulting from a peer production effort.

Security production is like the production of any complex good or service. It is not the result of a single process, but is the sum of several overlapping, parallel or complementary steps. The division of labour and specialisation in Internet security production requires a network of complementary or even intertwining production, in which one product is the resource for another product. Limitations of the openness of the result of one production process are at the same time a limitation of the openness of the input resources for another item. Openness in FLOSS is said to be viral, and so are the limitations of openness. Limiting the openness of one good or component in the security domain results in limited viability of peer production for another good.

The second aspect of the impact of secrecy on peer production refers to the consequences of non-openness during the production process and the production platform. One of the basic principles of peer production is that anyone willing to take part in the collaborative endeavour is welcomed to do so. But literature on open source projects tells us that closure is known to be a common organisational element of peer production. Many open source projects follow an onion model with layered rights. At the outer ring of the onion model, where the occasional contributors gather, anyone has the right to contribute and has access to all the technical and informational resources to be able to do so (Crowston & Howison, 2005). The right to include code into a release chunk of an open source software repository is, however, often limited to a rather small number of persons. With these hierarchical governance institutions, the process of building thus is not open for every contributor. Contributing to and using (intermediary) product however is. In larger open source projects, some functions in the production process are accessible only to the elites of these communities. But secrecy of the production platform imposes more limitations than hierarchical governance mechanisms. Non-openness/secrecy can mean the exclusion of potentially beneficial contributors to the entire production platform. This potentially runs against some of the economics of the viability of the peer production platform. One crucial element of peer production is that the set of contributing persons is not defined ex-ante by some kind of organising body, but rather is the result of distributed individuals' decisions to contribute or not. Only by this voluntary decision-making, peer production platforms manage to get the talent in numbers that make peer production a viable model.

Limiting access to the production platform requires organisational or technical barriers. The effectiveness of peer production suffers from secrecy and non-openness as vetting procedures and access-control mechanisms need to be implemented. These mechanisms increase the costs of maintaining the production platform. The motivation of potential contributors is diminished by increased transaction costs for contributors, caused by the introduction of vetting procedures. Some individuals may be willing to contribute, but they are not interested enough to undergo a process in which their trustworthiness is checked.

As a third impact vector, secrecy in peer production could result in reduced availability of the results of the collaborative efforts. As a defining criterion, the actual products that follow the peer production process need to be accessible, unrestricted and modifiable. Accordingly, non-openness can reduce the availability, the manner of usage of the product, and the ability to alter or reuse it. Limiting its openness in the sense of excluding others from using it, makes a good exclusive, and thereby alters the type of the good into either a private good or a club good. Security as a private or club good would severely reduce the legitimacy of the production process. This holds true in particular when there is little public involvement in the overall security governance. Exclusivity of security would reduce the legitimacy of the respective security governance arrangement. This would be in stark contrast to open source software that, due to their GPL licences, have the characteristics of a true public good. A second possible form of reducing the openness of the result of a collaborative production endeavour would be to limit its use and reuse. Given the virulence of non-openness as described earlier in this section, limiting the right to alter a good reduces its application in other collaborative production projects. This effect is independent of the question of the exclusivity of the good. The classification systems for information have greater variability than the binary exclusive/non-exclusive scheme of economist good classification.

Beyond the three dimensions of limitation of secrecy on peer production, there are several cross-dimensional effects of secrecy. First, product quality might decrease compared to complete openness as the flow of information on weaknesses and quality issues is hampered by secrecy measures. In addition, the barrier against malevolent actors at the same time cumbers the influx of potentially skilful resources. Second, secrecy can demotivate contributors. Peer production relies on the intrinsic motivation of contributors. The introduction of barriers between an individual contribution and its appearance in the final product implies the risk of demotivating contributors by excluding their contributions to final builds. These projects have thus created governance mechanisms to ensure that contributions are included completely, and in a timely fashion. The instruments of voice, exit and fork, the usage of which are free to anyone, create substantial incentives on the building actor to refrain from discriminating among contributors. Secrecy of the production platform and the products reduces public popularity and exposure of these communities, which also reduces a defector's ability to raise voice outside of the community and the general public.

2.3.4 Adapting peer production for security?

The previous subsection discussed various ways in which secrecy could impact upon defining characteristics of peer production in theory. What will happen in practice depends on actors' choices and their evaluation of the respective costs and benefits of secrecy and openness. If secrecy would affect all the elements of openness in the collaborative production process, it would look substantially different from the familiar FLOSS or Wikipedia communities. However, there might be ways to negotiate the tension between secrecy and peer production. Possible solutions could address any of the detrimental effects of introducing secrecy into a peer production and also address the underlying causations. The goal here is neither to build scenarios for all possible impacts of secrecy upon the entire collaborative production process nor to develop theoretical ways to mitigate the impact of secrecy. The remainder of this subsection seeks to showcase the fact that that adding elements of secrecy does not necessarily establish a production systems that is entirely different from peer production, and that, despite some limitations to openness, some characteristics could remain.

With accessibility restrictions to the production platform in place, the number of potential contributors is limited, but not as inevitably as it appears to be at first sight. Even if access is restricted, a layered approach or a modified onion model could still grant open access to a limited set of information at the basic level. However, any potential contributor would have to meet certain characteristics to be granted access to informational repositories. There are two aspects regarding why closure still has potentially detrimental effects. First, as mentioned earlier, any vetting or application procedure might deter potential contributors, and there is no way to avoid this but by lowering the entrance bar and making it seamless. This could either happen by some kind of automation or, probably more realistically, by introducing several layers of access. The second aspect, however, is impossible to avoid. Restricting access precludes the definition of characteristics by which the potential contributor is judged. This ex-ante closure impedes potentially valuable contributions by actors who do not live up to the norms defined in advance. A possible solution to the negative quantitative effect of closure on potential contributors might lie in multiple access layers, in liberal and flexible norms.

A further approach to address the negative quantitative effect of closure on potential contributors is to discuss the core assumptions of peer production theory. Peer production is a viable production model as: social value is created in quantities and qualities unmatched by other modes of production, superior knowledge creation is enabled by free access to information and by contribution of persons from different provenance, and they are then free to experiment due to the lack of managerial, hierarchical specifications. While peer production has a strong normative societal impact, it also manages to deliver results. Simplified, the assumed superior or at least competitive product quality is the result of a large number of voluntary contributors combined with a high degree of openness of information and low investment costs for contributors. From the angle of the end result of the production process, the number of potential contributors is not the decisive variable. The key criterion is the number of those actually contributing and the value of their respective contributions. A community that needs to close itself off could invent governance mechanisms and techniques that enable it to accurately identify valuable potential contributors. Simplified, the product quality would then be the result of the assumed likelihood of contribution, the number of potential contributors and the value of individual contributions. Hence, likelihood and value can make up for less potential contributors. If a community thus manages to bring in talents willing to contribute, if it is opened to a substantial number of potential contributors, it might be able to compensate for the losses of excluding a large number of potential contributors.

So much for the theorising about collaborative, distributed production under secrecy. The degree of secrecy required in reality, the impact on the production process, possibly governance techniques to mitigate the negative effects of secrecy on effectivity and costs of production — all of that might look entirely different from that which has been envisaged here. An exploratory journey into the world of Internet security production is to shed more light on these questions.

2.4 Peer production of Internet security

The rise of the Internet has nurtured a fascinating trend towards geographically distributed, self-organised collaboration driven mainly by individual persons. The poster-child of this new type of organisation in the production of informational goods has been free and open source software projects. Peer production and open source theories assume that this mode of production is generalizable and applicable to the production of goods other than software or encyclopaedia articles. The limits of the applicability of peer production though have not been thoroughly explored; the organisational results that might come out of mixing elements of peer production with hierarchical or market elements are not yet clear. This study tries to explore the boundaries of peer production by analysing the distributed collaborative production of Internet security.

The goal of this chapter has been to comb the existing literature for possible answers to the research questions, discuss the concepts necessary for this undertaking and build the theoretical foundation for studying the existence of peer production in Internet security, while exploring the impact on the production model by an environment not optimally suited for peer production. Secrecy — a feature that appears to be on the other, dark side of the security coin — by definition clashes with the openness necessary for the commons-based peer production-flavour of social collaboration. The previous sections of this chapter have explored first the ideas and theories of peer production and related forms of distributed collaboration. Furthermore, they have analysed the problems of Internet security, the governance challenges and approaches that have been taken so far. The subsequent section then looked at the role of secrecy and related it to the other core concepts of this study, openness and peer production. Secrecy, this study hypothesises, apparently contradicts the openness requirements of commons-based peer production and therefore probably is a major driver of the organisational design of Internet security production.

This section now seeks to join these bodies of literature, and define and discuss the concept of peer production of Internet security, its core characteristics and its theoretical feasibility. The first subsection summarises existing research that has already touched upon aspects of this concept. The second then clarifies the concept based on the finding of the previous sections of this chapter.

2.4.1 A brief literature review

The literature on what could be called the peer production of Internet security is rather scarce. Benkler (2004a) had probably the first shot with his article "Peer production of survivable critical infrastructures" presented at a conference in 2004. Therein, he discusses the viability and possible ways for "provisioning survivable critical communications and computation infrastructures by deploying radically distributed, peer-based systems for communication, computation, and data storage and retrieval." The concept builds on "excess capacity" of ICT services such as computation, data storage, and network connectivity. Using the willingness of users to share redundant capacities of their ICT systems and their subcomponents like disk space, CPU capacities, these "lumpy goods" as Benkler called them in a preceding article (2004b), resilient and secure computation, data, and connectivity services could be built using "peer-based survivable systems". Some of the ideas outlined in the article were implemented later on. Wireless mash-up networks using excess bandwidth capacities have been picked up in various open or shared-WLAN projects. The focus of these projects has however not been so much on security, but rather cost saving and increasing the general availability of wireless connectivity. In an earlier phase, the commercial backup service Wuala used a distributed backup model based on P2P-technology, allowing users to back up to other systems if they in turn granted some space of their hard-disk for backups by other users.

Johnson, Crawford, and Palfrey discussed ideas of "peer production of governance" for the establishment of an "increased online social order", which would help "address the collective action problems that arise in the online context." (Johnson et al., 2004, p. 7) The Internet would be plagued by various shortcomings like spam, network intrusion, or viruses that required regulatory interventions, best executed by "decentralized decision-making by establishing trust-based connections" (2004, p. 7). The authors apparently seek to avoid the rise of an "online dictator" that would enforce certain security mechanisms or "log all traffic" on the Internet to "know its citizens" (2004, p. 12). Their idea of "peer production of security" would in contrast delegate the authority over connections to the "individual level", introducing a "form of allowing access to send (non-filtered) messages dependent on (1) a demonstration of verifiable identity, and (2) the establishment of a basis on which the recipient should take the risk of allowing messages from the source to be received." (2004, p. 17) At the root of Internet security problems lies the "peculiarly unconstrained connectivity of the internet" (2004, p. 30). The social-technical order envisaged by Johnson et al. represents a significant departure from the then existing Internet architecture, but so is the since then established log-all-traffic approach. The authors sketch an approach to increase Internet security using peer production of governance. The paper understandably lacks what arguably requires a dedicated research programme: guidance of technical implementation of such trust-based connectivity systems and an analysis of its potential points of failures.

In 2006, then PhD student Max Loubser (2006) presented a brief research proposal titled "Peer Production of Security Information". Alas, the proposal never led to publications on that topic. Loubser's plan was to use social navigation to enhance the security of end-user machines and avoid the execution of harmful software. Social navigation, "the process of acquiring awareness and making decisions based on information gathered from others in a social space", would "leverage the community's collective wisdom to enable independent decision-making by the user about which programs to allow to execute." Apparently, Loubser cancelled this research plan and went for a more conventional topic. In 2011, he submitted his doctoral thesis titled "Organisational mechanisms in peer production — The Case of Wikipedia".[30]

A year earlier, L. Jean Camp and Allan Friedman studied "Peer production of Privacy and Security Information" as an approach to security production in computer science. The authors propose a system of security information that is not built upon a centralised data repository, but on a distributed, peer-based information exchange and patch mechanism. Their "peer-to-peer patch mechanism" was designed to counter that type of malware that was prevalent in the early 2000s, namely "good worms" (Camp & Friedman, 2005).

In 2009, William E. Koszarek wrote his master's thesis at the Naval Postgraduate School in Monterey on "Peer production in the US Navy". He has looked for clues as to how *Coase's Penguin* could be "enlisted", how the economic advantages of peer production could be reaped by the Department of Defense (DoD). The author has designed a rather generic template for implementing "peer production" processes in the DoD for maintenance, logistics, R&D and other internal, organisational functions (Koszarek 2009, p. 119). Possible volunteers for DoD-initiated projects could be retirees, college students, former professionals with lots of spare time and "the perpetually under-utilised or under-fulfilled" (2009, p. 87). While Koszarek offers some interesting ideas about the instrumentalisation of proprietary and social crowdsourcing, and peer-production — it is all labelled peer production by Koszarek — his work is somewhat underdone. More relevant here however is that his thesis contributes little to the question of security peer production even though he frequently refers to the US Department of Defence as a locus to apply his thinking.

These three papers are possibly the only attempts to link and conceptualise Internet security and peer production. There certainly are related bodies of literatures that touch on forms of social production of Internet security. The small literature on security communities is probably the most relevant among them. The notion of *community* resembles the organisational cell of open source production, where every substantial open source project seems to have its own community of contributors and collaborators. Brian Krebs (2006) was probably the first to put the spotlight

[30] Sparse details on the thesis are available at the website of the University of Oxford, http://ethos.bl.uk/OrderDetails.do?uin=uk.bl.ethos.522764.

on voluntary security groups. Mark Bowden's Conficker article (2010) and — later during the conduct of this study — even more so his book *Worm* (2011) share excellent accounts of the Conficker botnet and the response to it. With his journalistic background, he does not explicitly analyse the characteristics of Internet security communities. On the scientific front, networks of technological experts have been casually remarked upon in several articles. Van Eeten and Bauer (2008b) have based their "Economics of malware" study on interviews with "security experts" and "the security community". Discussing security governance on the Internet, Mueller (2010) provided a first glimpse into the world of "interpersonal and interorganisational networks" that seek to address Internet security issues like spam and phishing. Mueller, Ku, and Schmidt (2009) analysed the role and contributions of a specific community, the Anti-Phishing Working Group (APWG), for countering the phishing problem. Alexander Klimburg (2011) has recommended the use of "the large volunteer community" as a resource for national cyber-defence capabilities. Rain Ottis (2010) has published on the subject of the other end of volunteer communities, namely those more or less skilled computer users engaged in patriotic hacking or other forms of cyber attacks. Eli Dourado (2012) has analysed, apparently based on a literature review, "how ISPs enforce security norms". An in-depth analysis of the cooperation among operational experts in the NANOG community has been conducted by Ashwin Mathew (2010).

2.4.2 The concept of peer production of Internet security

The idea of peer production of Internet security is that security can be produced by joint efforts among peers. Peer production of Internet security resembles modes of production applied by open-source based software endeavours. It applies the organisational mode of peer production to the domain of Internet security. The organisational form, which is applied to reduce damages (and their likelihood) to values related to the Internet, incorporates the defining characteristics of distributiveness, openness, and socialness. In the narrow sense, the term is an abbreviation of the concept of commons-based peer production, which includes all three of the aforementioned characteristics. In the wider sense, peer production of Internet security only follows the criteria of distributiveness and socialness (cp. Section 2.1.4, Table 2.1). Embracing peer production as a mode of security provisioning and decomposing the security provisioning processes further extends the idea of non-state security provisioning that has been observed by policing and international security scholars (cf. section 2.2.1). The difference between peer production and these other trends in security governance though is that peer production does not rely on firms in the market, but is based on the principles of social production.

These characteristics of peer production of Internet security can manifest themselves at various stages or parts of the production model. The model was outlined in section 2.1.4 and visually summarized in Figure 2.1. Peer production of Internet security can refer to any process, activity, subprocess, and creation of intermediary products, supporting products, process-specific tools, and eventually Internet security conceived as a product (more on that later in this section). Most decisive is the actual production process as it includes many of the occasions where instances of peer production can happen. The list of elements of the production process has so far been generic. It can be supplemented by security-specific subprocesses. Professional literature on information security and, more specifically, incident response, often splits up the incident response process into several subprocesses (Cazemier, Overbeek, & Peters, 1999/2004, 2010; Clinch 2009).The purpose of such models of security processes is both to better understand and logically perceive the dimensions of security processes, and to optionally design organisational responses along these models. These sub-processes can include monitoring, incident detection, problem analysis, solution or mitigation design, implementation, or management and governance. In reality, in actual responses to Internet security incidents, these subprocesses may play a substantial role or not, and the activities and thinking of contributors might be shaped by such process models. The model of processes and subprocesses might help to structure information in initial phases of the empirical analysis.

What this social-scientific definition says is: a) peer production is an organisational approach that has the three aforementioned characteristics, b) Internet security is the low risk of damages to values that an actor does not want to see damaged, c) Internet security therefore can mean numerous things and cannot be defined irrespective of an actor's values, d) these values, which an actor wants to see protected, can also include the means to protect them, e) security production, i.e., reducing the risks, can include numerous activities and technologies. In the case studies, I have used a narrower aspect of Internet security, namely the necessity to re-establish the Internet's functionality after a large-scale incident.

The empirical part of this study will not focus on the institutions of Internet governance in the narrow sense, i.e., upon those institutions and organisations governing the core technical components of the Internet. Given the political forces currently putting pressure upon the Internet's components and institutions, it seems unlikely that future Internet security governance will be dominated by organisational modes like peer production, let alone in its commons-based variant. Nevertheless, these modes might have played and might again play a major role in the provisioning of Internet security. Before that can happen, however, some core problems regarding the applicability of peer production in security settings must be thoroughly understood.

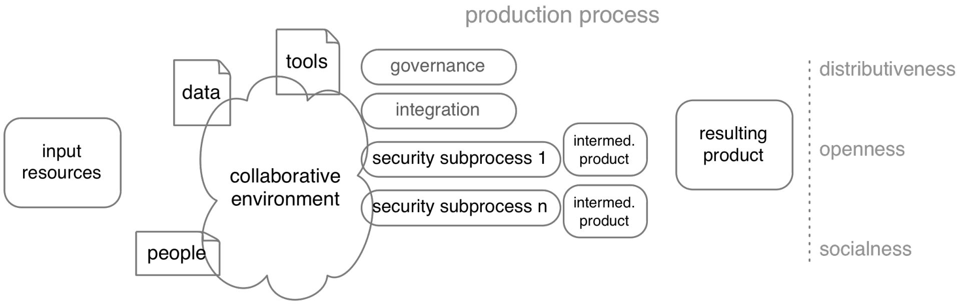

Figure 2.3: A model of the open source/peer production process and its security-specific subactivities

The discussions on openness and secrecy in the prior section 2.3 drew an ambiguous picture with regard to how much secrecy is actually needed, for what purposes, and which information objects it is actually applied to in Internet security production endeavours. In theory, secrecy could affect security production at various ends, as subsection 2.3.3 has shown. The openness of input resources could be reduced, the production platform could be restricted to members, or the resulting product could become a club or private good. What actually happens in reality is depicted and analysed in the later empirical chapters.

A few additional conceptual issues need to be cleared up before we proceed to the next chapter. Conceptualisations should help to achieve a common view on the topic of peer-produced Internet security. What has not been addressed so far are the differences between the notions of *security production* and *security governance*. Indeed, the term *security production* might raise doubts as to whether security is producible at all.

Security governance can be differentiated from security production in the following way. Governance can be described as the manner in which groups manage their common affairs, entailing an authority relation between those who govern and those subject to governance (Kahler 2003/2004). The activities of technical experts, who aim at stabilising the technical dimension of Internet security, ultimately result in altered code or configurations and hence change the way the Internet as a technical infrastructure can be and is used. Therefore, their activities towards providing Internet security can also be described as Internet security governance.

Some claim that security is not a *product*, and could therefore not be *produced*.[31] This reservation is best addressed by referring to fundamental definitional characteristics of the term security. Security is a multifaceted concept. On the one hand, it is one of the fundamental value ideas of modern societies, a normative guide which strives for an unattainable state of mind and world, in which control over the future and prevention against risks and threats are guaranteed (Kaufmann 1973). Security in that sense would be somewhat difficult, if not impossible to operationalize. However, there is another meaning in "security" that refers to security as a man-made quality (1973). If security is something producible by human action, then it can have the qualities of a good or a service. As such, it can either be a public good or proprietary and owned (Krahmann 2007, 2008). Furthermore, policing studies have yielded a literature that focuses on security as a producible good. Security here is the result of different processes, sub-processes and tasks pursued by different actors, no matter if private or statist (Caparini 2007; Hänggi 2003; Kempa et al., 1999; Bowling & Newburn, 2006; Bryden & Caparini, 2007).

[31] Security professionals often argue along these lines to stress that information security is not the result of buying a range of security products from some security vendors, but instead continuous processes combining tools, knowledge, and certain operational activities.

3 Research Design

The rise of networked governance and networked collaboration has created an interesting puzzle: What happens, which modes of governance and production arise, when the feasibility criteria of peer production are mostly, but not entirely met? The problem that this research endeavour aims at addressing has been described in chapter 1 and elaborated on in chapter 2. We lack knowledge about what happens to the idea of peer production and its applicability as a mode of production when the circumstances are not entirely in its favour. The current literature on open source and peer production has little explanatory, let alone predictive, power as to whether the peer production model is applicable in inconvenient environments. Nor can it easily explain which hybrid forms of production might emerge if some characteristics of peer production are merged with characteristics of traditional forms of production and governance. Furthermore, the overall societal relevance of peer production as a different form of production is far from clear. While some authors have evangelised the game-changing nature of peer and open source production, neither the quantity nor the quality of applications of peer production have been systematically evaluated. It is not clear how common a phenomenon peer production actually is, how applicable it is to the production of goods other than software or encyclopaedias. While substantive research on the open source method and peer production has already been conducted in the domain of software production, empirical research only slowly progresses to other societal domains. Security is one such domain. It arguably is an inconvenient area for peer production because of potential secrecy requirements, and security arguably is the branch of politics in which, to use a common definition of politics, authoritative and generally binding decisions are made and literally enforced.

The current state of literature on open source and peer production only allows for deductive and interpretational answers about the applicability of peer production. The literature tells us about possible hindrances to the applicability of peer production. The previous chapter has contributed additional thinking about theoretical hindrances to peer production in the domain of Internet security. This hypothesising has itself been based on hypothetical, theoretical thinking about the feasibility of peer production. The objective of this study is to analyse organisational approaches to Internet security provisioning and thereby help to understand the theoretical and empirical limits of the applicability of peer production. The research questions follow this goal. The first question aims at summing up the organisational status quo in a particular segment of Internet security production. The subsequent questions seek to shed some light upon the reasons for the application and non-application of elements of peer production.

To contribute to the underlying theoretical debate and to answer its research questions, this study uses case studies and qualitative expert interviews as the main source of empirical data. The empirical analyses focus on two cases of large-scale incident response: the Estonian cyberattacks and the Conficker botnet. Gathering data in the domain of security politics can be tricky. This study, like all others, cannot avoid the need for compromise in the research design. The purpose of the two cases approach here is not so much to compare the two of them, but rather to extend the empirical base, to increase the validity and significance of possible outcomes. This study is an exploration in the empirically rather unexplored territory of Internet security production viewed through the lenses of the peer production model. Furthermore, it seeks to explore the role of the so-called Internet security community that has been mentioned, but not thoroughly studied in previous literature (cp. section 2.4.1). A case study on incident response will provide more details about the contributions of the communities in practice. A second case study broadens the empirical base and helps unveil the characteristics of the community in more generic terms and different approaches in different geographies.

This chapter is organised as follows. The next section elaborates upon the overall research strategy. It describes and explains the major design decisions taken during the course of this research project, among them the criteria for the case selection and a brief expose of the two cases. The following section summarises the models on identifying peer production and explaining its applicability and hindrances in the domain of Internet security. The last section describes the methods for collecting and analysing data.

3.1 Research strategy

3.1.1 Incident-based case studies

Internet security incidents are events or developments that put at risk that which actors have defined as pivotal to their Internet-related interests. These events challenge securing actors and their abilities and procedures to adequately deal with these risks and incidents. By creating a stressful situation for the actors involved, they expose the strengths and weaknesses of existing institutional mechanisms and workflows at a given moment of time. They highlight existing modes of providing security beyond and independently from what has been planned in policy papers. Action and non-action, resources, processes, modes of production, and outcomes are what matters in incident response. The study of incident-based cases hence allows for a robust description of actors involved in security provisioning, the organisation of their activities, their access to relevant information, interactions and collaboration among several actors and their operational problems when dealing with demands for secrecy and the need to exclude certain actors.

The incident-based case studies approach comes with some methodological implications. First, the analysis of security incidents will not reveal all modes of security production or every aspect of security provisioning. This analytical perspective focuses on the direct activities of the persons and organisations involved and therefore does not include a fully fledged analysis of all preventive measures. Hence, it excludes a substantial branch of security strategies from the empirical analysis. If prevention is about reducing opportunities for perpetrators or non-actor based technical vulnerabilities, incident reaction analysis looks at what happens once prevention has failed and risks have turned into immediate threats for security. Hence, this study will not assess modes of production used for preventive measures – except, of course, for the resources and procedures that have been set in place for incident response. Secondly, the analysis of security incident responses implicitly supports certain conceptualisations of Internet security: the one that actors involved have defined. An incident response approach hence perpetuates given understandings of Internet security no matter how much these conceptualisations are flawed from a social science perspective. The incident response approach does not provide an analysis of the exploitation of biased functional concepts of Internet security. A case selection that focuses on security incidents that do affect the technical functionality of the infrastructure itself can mitigate this problem. The technical foundations of the Internet should be less politically contested than incidents relating to specific usages of the Internet. It would certainly have been interesting to look at how the contestedness of security definitions and the effects of fundamental political or normative disputes among potential collaborators affect openness and the willingness to use secrecy for self-interest purposes. Wikileaks and the Snowden leaks have perfectly illuminated the scope of potential security incidents. Third, an incident response approach based on distinct, independent cases with only little distance of time between them is not perfectly suited to provide insights into diachronic developments. Changes in the way similar incidents are handled over time will not be revealed, which is of course a pity for everyone interested in historical developments and ongoing changes in Internet security governance. This holds especially true as the political landscape for Internet security has significantly changed after the cases analysed. However, this thesis focuses more narrowly on gathering data adequate for answering the research questions. As one of the research questions asks how the balance between openness and secrecy is handled in a particular case, a synchronic approach is more promising than one in which changes over time are examined.

An alternative approach would have been a comparative longitudinal study, e.g., by comparing the approaches to similar incidents over time. A longitudinal analysis of botnet response would have been a viable approach. It might have yielded insights into the changing regulatory landscape, new institutions, and the role of corporate programmes for incident responses. But the focus here is to explore the application and viability of peer production of Internet security, and not so much the organisational alterations of Internet security production over time and an up-to-date depiction of Internet security governance.

Despite all these caveats and methodological implications, an explorative incident response approach seems appropriate for analysing Internet security governance. The overall governance structure and response capabilities in the security area are generally driven by incidents. Other than imaginative worst-case scenarios, incidents and the response to them show the essence of actual Internet security production and which resources, actors and knowledge are important for a 'secure' Internet.

3.1.2 Case selection

Internet security problems can come in many flavours, depending on the objects that are endangered, the technologies used, the actors involved. There are different kinds of Internet security problems such as the endangerment of technical IT-security or general technical functionality of the Internet, compromise of the functionality of infrastructures built on top of it, degradation of Internet-based business activities, or the financial risks caused by Internet-facilitated theft.

A number of further potential cases were pondered, but eventually quashed. Among them are cases related to online content, highly classified national security-related incidents, or the analysis of software communities creating software and tools used to respond to incidents. Certain forms of Internet content, ranging from theft of copyrighted content, child abuse material or the publication of terrorists' manuals is regarded as a threat to the security interests of some actors. These cases however refer to dimensions of security above the infrastructural-technological layer (used in the sense of OSI layers). They are linked to value- and interest-based considerations and definitions of security that are not related to technical security. Such aspects of security do not meet the criteria of a narrower concept of Internet security as re-established functionality after an attack as defined at the end of section 2.4.2.

This also holds true for cases that have been dealt with by secretive national security institutions. For such cases, data collection would have been a major burden and possibly a deal breaker for the viability of this research project. This however does not exclude the possibility of considering the relations between traditional national security institutions and their collaboration with Internet security communities in incident response.

Another approach would have been to analyse the peer production of security resources and tools. An equivalent to this resource-oriented approach would have been the analysis of, e.g., Whois or phishtank.org as distinct cases. However, the production of Internet security involves more than resources, it also includes actors and their modes of collaboration. Internet security provisioning can, at least theoretically, make use of modes of peer production with regard to incident response in two ways, a) by using open-source produced tools and, less narrowly, b) by collaboratively producing information and resources, their sharing and common access to achieve what could be called Internet security. To increase the relevance of the findings of this study, the second path was chosen. An interesting phenomenon in recent years has been the rise of distributed, collaborative research on Internet-based attacks on computer networks. One example of similar open source intelligence endeavours is the so-called *Project Grey Goose*, the collaborative, voluntary effort of a number of Western security experts who studied Russian web fora after Georgian Internet-connected systems were brought down by a series of Internet-based attacks (Project Grey Goose 2008). This and other projects, however, did not include an active response beyond the mere analysis, and were therefore dropped. Furthermore, Project Grey Goose has been transformed into an actual for-profit consultancy after the publication of their first report (Project Grey Goose 2009).

The research questions focus on the relationship between Internet security, openness and secrecy. This requires studying incidents that are limited in time and scope. These cases were selected by criteria such as the significance of an incident and the availability and accessibility of data on the incident and the response measures taken. In addition, the answer to the research questions should not be trivially apparent for a case. Many cases especially in the domain of military cyber-espionage have been dealt with by traditional security organisation. For such cases, the answer to the question of the application of peer production would likely have been a trivial 'no'; in addition, a few brief encounters with military CERTs from European countries convinced me that getting research data out of them would likely be a futile endeavour. Hence, the institutional approach toward the security problems should include some elements of peer production. The case study approach here needs to reveal whether responses to incidents are peer produced and whether these responses build on peer-produced resources.

Eventually, the Estonian cyberattacks and the Conficker botnet were chosen as the cases for empirical research. The cases appear to differ in a variety of ways: the technologies used for attacks and defences, the actors doing the attacking, the affected and the security-providing sides, the scope and scale of the problems, the degree of national security interests affected by the attacks. More importantly, however, they similarly match many of the criteria defined for the case selection. They both represent large-scale, international incidents; there are significant indications of distributed, informal collaboration, and indications of security provisioning not driven by markets and hierarchies.

The Conficker botnet first appeared in November 2008 and since then has risen to a network consisting of millions of infected computers. As the botnet has so far not been used for criminal activities, the actual damages caused by this botnet are limited to the costs of botnet mitigation, machine disinfection and the like. However, the combination of advanced malware techniques, cutting-edge cryptography and peer-to-peer technology has made this botnet extremely difficult to counter. The response has attracted contributors and spurred collaboration among Internet security experts worldwide.

Botnets are technical platforms that create a number of technical, economic and political risks for the Internet. A multitude of client machines, so called bots or zombies, are infected with malicious software that is remotely installed on those machines exploiting vulnerabilities within operating systems and applications. These bots can then be used for malevolent purposes, orchestrated by one of the several machines that make the command and control layer of a botnet. Anti-botnet measures generally can encompass a wide range of activities such as problem analysis, analysis of the design of a botnet, development of mitigation techniques and software, monitoring of botnet activities, implementation of mitigating techniques of botnet activities, destruction of botnets. The investigation of this case will determine the degree to which the response to the botnet, including actors in the so-called "Conficker cabal", relied on peer production methods, and will also analyse the way it handled the problem of secrecy and openness.

Between the end of April 2007 and the middle of May, a series of cyber attacks targeted Internet-connected ICT systems of Estonian banks, media corporations and ministries. Their web-based services were in serious trouble and at times even entirely unavailable after a countless stream of distributed denial of service attacks (DDoS), defacements and other forms of attacks.[32] According to somewhat hysterical members of the press, a country was brought down by cyberattacks.

DDoS attacks are attempts to interrupt the availability or decrease the efficiency of computational resources by using large numbers of coordinated attacking machines to flood the target machines with high numbers of requests beyond their capacities. The discovery of abnormalities, the analysis of the causes of the Estonian events, and mitigation activities were based on distributed global collaboration. The investigation of this case will determine the degree to which actors responding to the Estonian DDoS attacks relied on peer production methods and the way they handle the problem of secrecy and openness.

3.2 Analytical model

An analytical model tailored to the research questions of this study was outlined in the previous theoretical chapter. It follows the main research questions that ask, first, for the existence of peer production or elements thereof and, second, for possible reasons for the non-application of peer production, especially regarding the role of secrecy and the presence of antagonists. The empirical research consequently consists of two parts. The first part encompasses a description of the production process, an identification of the products produced and the most relevant required and used during the production process. The second part identifies and discusses the factors that hinder or promote peer production.

[32] For a brief technical analysis cf. (Nazario 2007), for political implications cf.(Traynor 2007; Anderson 2007).

3.2.1 Identifying peer production in incident response

Identifying the application of peer production in incident response requires three steps. The first task is to identify the goods and services produced by the response activities. This task requires a solid narrative of the actual response activities and the identification of the most relevant elements of the response. The second task is to classify the response and its most important components, using the specified criteria of peer production, namely distributiveness, openness, and socialness. The third task is to decide whether the response as a whole can be classified as peer production or not.

As the first step, the response endeavours are analysed and their respective outcomes identified. Section 2.4.2 has provided a framework for characterising Internet security production. It breaks Internet security production down into several *elements* including input resources, intermediary results, and outputs of several security processes, among them monitoring, analysis, mitigation, forensics, sanctioning, resource allocation and governance. Using this framework, the details of the actual production of Internet security are to be analysed. Actual response activities might be characterised by processes that are more specific or by an organisation with a less elaborate division of labour. On the one hand, this requires an analysis of the actual production of Internet security in the most meticulous and detailed manner. On the other hand, only decisive and crucial parts of the response activity are of interest. Discussing the peer-producedness of marginal parts of Internet security production would yield only marginally relevant results. As an example, the use of open source commodity software is irrelevant for the underlying questions of the relevance of elements of peer production in Internet security. Therefore, the execution of the case studies needs to balance depth, coverage and relevance.

The second task is to analyse the response through the lenses of the peer production model. The model of peer production of Internet security described in the theoretical chapter helps to identify elements of peer production in the actual responses in the two cases. This model then needs to be applied to the overall production process and its subcomponents. A collaborative production system can be called peer production if it matches the criteria of *distributiveness*, *openness* and *socialness*: Collaboration among participating actors is distributed or highly decentralised;[33] resources and information required during the production process are shared ad-hoc and in an unrestricted way, and produced goods can be reused and

[33] Cp. footnote 112 for the blurry lines between distributiveness and decentrality.

adapted by anybody involved in the production system; the actual production processes are not driven by market incentives or hierarchical commands.

As a third step, the most important activities, processes, intermediary products, products, and eventually the entire response need to be categorised along the three defining characteristics of peer production. The question as to whether the response to an actual Internet security incident can be called peer production, entails a qualitative and a quantitative dimension. On the qualitative dimension is the importance and relevance of the response elements and their peer-producedness. It is possible that the bulk of the security provisioning processes and goods are the result of market-based or hierarchically organised activities and transactions, while some crucial resources or intermediary products are peer produced. In such a case, the overall response might not be peer produced, but peer production would still have a viable and significant role. The relevance and importance of a particular response element will be determined by the assessments of the interviewees.

On the quantitative dimension is the *number* of elements of peer production found in the response activities. Internet security governance consists of several steering and provisioning processes (cf. the paragraphs on the process dimension below). It is conceivable that these elements will have most of the characteristics of peer production, but not all of them. If this is the case, we might observe the kind of innovation in production processes that Weber predicted would happen when modes of open source production are challenged by particular circumstances. In this case, the challenge would come from the need for secrecy in security production. Therefore two important quantitative aspects are likely to be present in each case: a) the number of sub-processes of security production that involve modes of peer production and b) the number of elements of peer production (distributiveness, openness and socialness) that are present throughout.

This study shall focus more on qualitative aspects. That design decision is a response to a problem caused by counting the quantity of peer-produced goods in security provisioning. It is conceivable that substantial or minor numbers of goods and services of Internet security provisioning are the result of peer production. Measuring these quantities would require an all-encompassing blow-by-blow study of the provisioning processes. Such quantitative analysis would have to allow for relative statements like an estimate of how much of a security problem is addressed by one mode of production and how much by another. Given the difficulties of the data collection process among a somewhat secretive community of experts, these questions might be impossible to answer consistently.

The classification of the elements of the response can result in four possible answers to the question of the peer-producedness of Internet security production. *Total peer production* and *no peer production* are conceivable results, but are unlikely to occur and can hence be disregarded. Given the initial desk research and current political trends of increasing state involvement, the impression at the outset of this research project was that it is highly unlikely to find pure peer-production systems for incident response activities. In *total peer production*, every activity of the overall incident handling could be subsumed under the label of peer production. With *no peer production*, no activities would make use of open data repositories or apply peer production-like forms of collaboration. It is more likely though that the response to an Internet security incident falls into one of the following two categories. *Substantial peer production* occurs when qualitatively important response elements are provided by peer production or by a response in which important elements approximate the defining characteristics of peer production. In a quantitative perspective, many elements of the Internet security provisioning process are provided by the peer production method. With *minor peer production*, only response elements of minor importance are peer produced, while important response elements only display some of the characteristics of peer production. Quantitatively, only a few elements of Internet security are provided by modes of peer production.

Table 3.1: Classifying the peer-producedness of security production

Total pp	All response elements fulfil the characteristics of pp.
Substantial pp	Pp'ed security elements are of importance for the response. Response follows many of the defining characteristics of pp.
Minor pp	Pp'ed elements are of minor importance. Response elements follow only some characteristics of pp.
No pp	None of the response elements even remotely resemble pp.

A few words on the organisation of this book and in which segments these tasks are addressed. The relevant response elements are described in Chapter 5, Producing Internet Security. The analysis of the response along the defining criteria of peer production and the classification of the response elements is the purpose of Chapter 6, Social Dimension of Internet Security.

3.2.2 Hindrances of peer production

The second set of research questions asks why elements of peer production were or were not used in responses to Internet security incidents. The model of the feasibility of peer production, the result of a literature survey and supplemental theoretical work, describes factors that have been identified as likely prerequisites for the feasibility of peer production for the creation of a particular good.

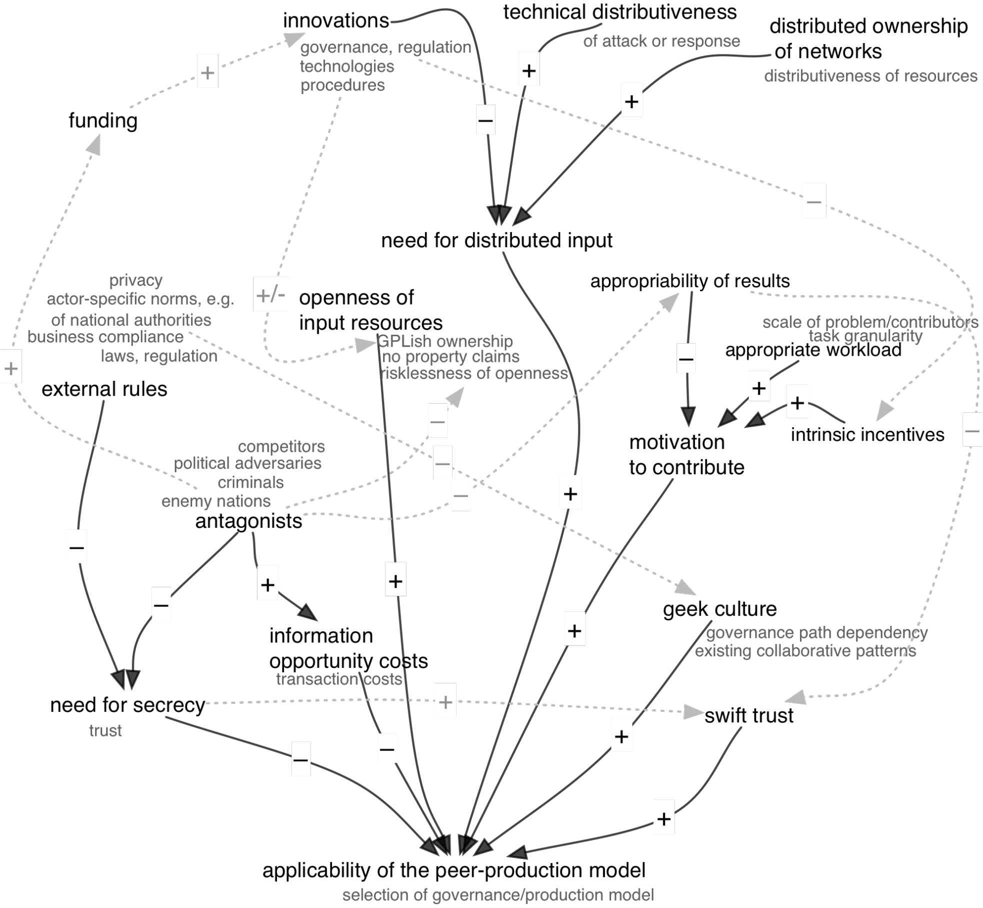

Figure 3.1: Drivers and hindrances of peer production

The result is an unsurprisingly complex figure of lines and arrows between a large number of concepts and subcategories like distribution of ownership of networks, distributiveness of attacked system and required response systems, need for distributed input, innovations, motivations to contribute, community and culture, information opportunity costs, openness of input resources, need for secrecy, external rule, funding plus a substantial number of subcategories for many of the enumerated concepts (cp. sections 2.1.3, 2.1.5, and 2.3).

The second set of research questions is particularly concerned with the role of secrecy and of the antagonists, and their influence upon the mode of production chosen for the response activities. The first sections of Chapter 7, Limits of Openness, develop a model of the drivers of secrecy. The antagonists are one such driver, as are state and national policies, businesses' economic interests, and social and regulatory norms. The role of these factors pushing toward secrecy in the particular response endeavours then needs to be asserted. Epistemologically, this does not and cannot result in the discovery of causal relations. The complexity of the conceptual relations and the interdependence of feasibility factors of peer production apparently cannot be reduced to a very small number of controllable variables if Internet security production in general and incident response in particular is the chosen empirical domain. Thus, the sort of knowledge derived from this part of the analysis will not reside in the league of 'proven causal relations'. This study is exploratory; it aims only at yielding well-researched assumptions, models and hypotheses on different relations of peer production and Internet security production.

3.3 Data collection and analysis

This section elaborates upon the methods used to gather the data required to answer the research questions in the manner outlined above. Among the data collection methods used are desk-based research, qualitative expert interviews, and a bit of direct observation. A description of the processing of this data follows, supplemented by a few comments on the means applied to increase the validity of the findings of this research.

Desk research was the chosen means of data collection for a variety of aspects of this study. Obviously, grasping the state of the literature on those knowledge domains touched on in this study required searching countless articles and a few books. Search terms like Internet security, Estonia 2007, Conficker, botnets, security production, security governance, secrecy, openness, communities, cyberwar, Internet governance, networks, organising security, security economics, international security, trust, information sharing, peer production, open source, social production, including variants or synonyms of these terms, have been entered into various Internet or science-oriented search engines. Furthermore, the bibliographies and footnotes of existing relevant literature were reviewed for more valuable writings; blogs of scholars and practitioners in the fields studied here watched; their increasingly numerous Twitter accounts followed and valuable links shared, followed and read. Found articles, chapters, books or reports have been placed into a document storage, organisation and analysis software and, if sufficiently relevant for an in-depth analysis and possible later citation, into a dedicated bibliographic management software. Inevitably, too many of the thousands of texts gathered during the course of this research project still have a "worth a read" tag attached. Next to its invaluable role in helping to get on top of current debates and grasping the state of literature, desk research has helped to cross-check some statements and make sense of the security community's special lingo heard during the interviews. The aim of the desk research has also been to understand the technical, organisational, economic and political details behind incidents and the responses to them.

The bulk of the data needed to answer the research questions underlying this study came from interviews. The literature on the empirics of Internet security production has been, and still is, far from overwhelming. The same holds true for the empirical cases of this study in particular. Published data on the incidents and the responses to the Estonian 2007 cyberattacks and the Conficker botnet are inadequate. Thus, additional data on the cases has been gathered from semi-structured qualitative interviews with actors involved in incident handling. Expert interviews have been the indicative research strategy as knowledge about the organisational issues of large-scale incident response and global collaboration has been scarce. The reluctance of large parts of the security community to share information with the public has prevented most knowledge about its inner workings from escaping into the public realm and become publicly accessible, explicit, written knowledge. The interviews were designed as semi-structured to ensure that they yielded sufficient knowledge on the different aspects of this research, while giving the interviewees sufficient leeway to share their own perspective and aspects that had not been envisaged in the design phase of this research project. The aim of the interviews has therefore been to learn about the actors involved, their motivations, the forms of governance used, the processes of security provisioning, the processes, tasks, and products required to adequately respond to the incident, the forms of interactions among technical experts, the role of secrecy and the impact of the antagonists on the organisational form in order to be able to answer the specific research questions. The interviews were conducted based on an interview guideline that included the different aspects of the research questions and the analytical models designed to the problem of identifying peer production and explaining its application or non-application and the impact of secrecy and the antagonists thereon.

Interviewees were selected according to their involvement in or first-hand knowledge of the incident and especially the response activities. I learned about the involvement of those persons based on press articles, and then later in the interview phase by snowballing, i.e., by hints from other interviewees. Another decisive criterion was the accessibility of the interviewees. A few candidates were not interested, a preference conveyed mostly by not replying to direct requests for interviews. Twenty-seven formal interviews were conducted, mostly between January 2011 and June 2012, with a few follow-ups later. They lasted roughly one hundred minutes on average and a good forty-three hours in total; the two short ones were finished after thirty minutes, the two longest ones almost covered a working day. The interviews were mostly conducted face-to-face, a few though via Skype with video, some audio-only or via telephone. In-person interviews were usually and necessarily and for better or worse conducted in bars, cars, cafés, hotel lobbies, restaurants, and other environments with loud music and booze in the course of conferences in Tallinn, Vienna, Rotterdam, Frankfurt, Washington, New York, San Diego, Los Angeles, and San Francisco — just as textbooks on qualitative research life for the very social sciences would recommend. While these environments might have had deteriorating effects on the mental sharpness of the interviewer, they might also have created the relaxed atmosphere that is necessary to talk about security issues. Almost all interviews were conducted in English, a few in German. Most in-person interviews were recorded. In almost any recorded interview, interviewees asked to pause the recording for a while to share some details off the record. Few interviews were conducted without any recording. For those interviews I took brief notes, usually only keywords, and wrote down memory minutes right after the conversation. During this research project, I attended a number of conferences and workshops on Internet security such as a FIRST conference, security sessions at various Internet Governance Forums, the Dutch GOVCERT conferences, and private community meetings. It is these security conferences where members of the security communities, who usually collaborate online, meet in person. Attending these meetings is certainly a far cry from methodologically conducted participatory observation, but it allowed me to get a better idea of how these communities work.

To analyse forty-three hours of interview material with underlying base beats, pelting rain, wheezing espresso machines, hubbub, and howling engines in a relaxed, content-focused manner, all recorded interviews were transcribed in a pragmatic way, i.e., in plain text without meta-information about tonality, pauses, and the like (Kuckartz 2010, pp. 40-50; Dresing & Pehl, 2011/2011). The transcripts and, for those interviews that had not been recorded, the summarising post-interview notes, were then coded with an established QDA software. Coding scheme mostly followed the previously described conceptual and analytical models and were designed to understand the incidents (among the sub-codes here: attack description, technologies, context, attribution, damages), Internet security production and the response to the incidents (activities and sub-processes, team formation, organisational issues, resources, impact, critique), identification of peer production (motivations, managerialness, outcomes, access restrictions, modes of collaboration, sharing), understanding hindrances and drivers of peer production (secrecy, antagonists, state policies, regulation). Some codes were added later or in-vitro like the security community and its sub-codes such as practices, values and views, trust and structure. The purpose of coding has foremost been to increase ease of access to statements on similar concepts and categories.

A variety of measures were used to ensure the validity of the findings of this study. Statements of interviewees were validated against those of their peers; facts heard in the whole body of interviews were cross-checked, if possible, with literature by other scholars or press articles. The analysis of the Estonian events has been published in a well-received edited book on the history of cyber attacks, after it had gone through a thorough peer-review by persons with intimate knowledge about the security community and the events back then (Schmidt 2013a). The narrative of the Conficker incident is largely consistent with Mark Bowden's account of the Conficker response, which was published in the midst of the data gathering phase for this book. The overall design and preliminary findings of this study have been presented in a peer-reviewed journal article (Schmidt 2012). Aspects of my research such as core concepts, research design and findings at various steps of my project were shared with the scientific community at several conferences hosted by the International Studies Association (ISA) or the Global Internet Governance Academic Network (Giganet). The implications of the findings for International relations have been discussed in a co-authored, peer-reviewed paper (Mueller, Schmidt, & Kuerbis, 2013). Essential findings on the security community and some thoughts on them that extend beyond the narrow question of this thesis were developed in chapters in edited books and partly also shared with the security community. Going beyond the mere research questions of this study, but addressing the underlying research problems and interests, a contribution to an edited book has considered the theoretical and empirical possibilities of hybrid form and cooperation between the networked community and traditional hierarchical security institutions (Schmidt 2014). The idea of social production, its diverse forms, and the Internet security community have been presented in another article (2013b).

Last but not least, two remarks on research ethics. The first touches on reusing previously published material mentioned above. Those texts, paragraphs or subsections that have been reused in this thesis are not specifically marked. Most of that material was created in drafts of sections of this thesis in the first place. The second remark refers to the identification or rather anonymity of interviewees. As a rule of thumb, interviewees remain pseudonymous in this work. In addition, I have sent excerpts to almost all interviewees and asked for comments and their preferences regarding anonymity. The excerpts contained paragraphs of this manuscript, in which they have been discussed or mentioned as a source. Some interviewees wanted to remain entirely anonymous; some asked not to be identifiable as a source for certain statements; some statements can possibly be traced to a person when only a few persons provided answers for certain questions.

4 Endangering the Internet

Every once in a while the Internet's core functionality, its ability to route information from one endpoint to another, is endangered by a new threat, a new botnet, a new attack that appears to have the capacity to bring it all down for a significant chunk of the Internet populace. Alarms are sounded by a few security specialists at first, then they are joined by others, until the tech press chimes in, blogs are filled with reports and eventually mainstream media conveys the news of the acutely critical situation into every household. The Estonian cyberattacks in 2007 and the Conficker botnet have been two such incidents.

Such large scale incidents challenge those involved with daily Internet operations, both individuals who have had a long-term interest in all matters Internet security and organisations that are responsible for certain artefacts that make up the Internet. Both incidents give a clear impression of how the Internet is "secured" and how it is "defended" in times of a crisis. It shows who uses which resources to create which products or services that eventually re-establish the *status quo ante* or at least a mitigated, stable situation, in which the Internet is again able to reliably fulfil its communicational tasks and functions.

This chapter gives a thorough description and narrative of both incidents, based on desk research and interviews with persons involved. The incidents, the DDoS attacks and the botnet, create the general set-up and stage on which the responding actors 'produce' Internet security, as later described in Chapter 5, Producing Internet Security.

4.1 Estonian cyberattacks: A country sort of under attack

For three weeks from April 27 until May 18, 2007,[34] components of the Estonian Internet infrastructure operated by governmental, political and administrative institutions, the private sector and some individuals were subjugated to Distributed Denial of Service (DDoS) attacks, website defacements, DNS server attacks, mass e-mail and comment spam. The attacks are supposed to be the first that were likely directed as an instrument of political conflict against a whole, albeit small country. At the time of the attacks, Estonia was locked in a domestic conflict between the newly elected government and its supporters on the one hand, and the Russian ethnic minority group on the other. In addition, the long-standing conflict with Estonia's former occupant power Russia culminated in heated diplomatic exchanges at a time when Russian-U.S. relations approached their post-Cold War bottom.

The incident is noteworthy for more than its geopolitical connotations. It also sheds light on organizational aspects of cybersecurity and the distributed collaborative nature to re-establish the Internet's functionality after an attack.

This section gives a descriptive account of the attacks and the damages it inflicted. This narrative is supplemented by an analysis of the political circumstances of the attacks, the discussions it spurred, and some recapitulating remarks.[35] The response to the attacks is then later described in section 5.1.

4.1.1 Monument debates

In January 2007, the Estonian government announced that it would move a World War II monument from the centre to a military cemetery in the outskirts of Tallinn. Erected in 1947 when major affairs in the Estonian Socialist Soviet Republic were controlled from Stalin’s Kremlin, the “Monument To the Fallen in the Second World War” depicts an unnamed soldier wearing a uniform of the Red Army, a helmet in his left hand, his head slightly bowed as if he would mourn his almost 11m fallen comrades ("World War II Casualties," 2012). After Estonia had regained full political sovereignty in 1991, the monument emerged as a point of conflict in domestic Estonian affairs. Many Estonians regarded the Bronze Soldier, which was located at a busy intersection close to Tallinn’s picturesque historic cen-

[34] At the end, the attacks frayed out a bit, hence the end is not as sharply delineated as the beginning. Therefore, in some descriptions May 23 is given as the end date and 3 ½ or 4 weeks as the overall duration.

[35] An earlier version of this chapter has been published (Schmidt 2013a).

tre, as symbol not of the achievements of the Red Army in WWII, but of its subsequent role as a suppressor of Estonian independence. Russian-Estonians begged to differ. Unsurprisingly, the monument emerged as the site where different interpretations of the role of the Red Army were expressed in demonstrations. The date of May 9, the Russian V-Day,[36] had become notorious for not only verbal clashes between Soviet war veterans and Estonian-Russians on the one side and conservative Estonians on the other. After years of repeating rallies, discussions about the future of the monument and demands for its removal gained momentum in 2006 (Alas 2006).

It didn't go unnoticed in Moscow that its former province in the Soviet era was about to cut cordons to the Russian interpretation of Estonian WWII and postwar history. The Kremlin was cross. In January, the Russian Upper-House filed a resolution demanding their Estonian parliamentary peers to halt legislation that would remove the monument. On April 3, Russian First Vice Prime Minister Sergei Ivanov made a plea for boycotting Estonian goods and services, though this bullying stance was not shared in Russia's foreign policy circles ("Here We Go Again," 2007). The conflict was about Estonian identity, relations between Russia and Estonia, and the perception of World War II (Myers 2007). For Russians, it was the Red Army that wrestled down the German war machinery in the bloody battles of the "Great Patriotic War", which cost the lives of approximately 27M USSR citizens (Kosachev 2007). In the eyes of many Estonians however, the Nazi occupation was only relieved by five-decades of Soviet occupation that continued the suppression of their striving for autonomy (Socor 2007).[37]

After having smouldered as a divisive and emotional issue in Estonian politics and public discourses, the monument eventually became one of the core subjects in the run-up to the Estonian parliamentary elections on March 4, 2007. "War graves are no place for day-to-day politics", warned President Toomas Hendrik Ilves, a Social democrat, but to no avail (Alas 2007b). The Union of Res Publica and Pro Patria, a conservative opposition party, lobbied for a bill prescribing the removal of the monument. Trailing in the polls, the incumbent Prime Minister Andrus Ansip and his Reform party supported the controversial bill in February, fearing an electoral

[36] The Allied Forces had summoned Wehrmacht General Jodl to Reims, France, on May 7, 1945, to sign the capitulation, to be effective on May 8, 23:01 CET, i.e., after midnight in Moscow. In addition, the Soviets held another signing ceremony in Berlin on May 8, close to midnight ("German Instrument of Surrender," 2012; "Victory Day (9 May)," 2012).

[37] For a more thorough discussion on the „memorial wars“ in Estonia cf. Pääbo (2008), for memory politics in Europe and May 9th, cf. Onken (2007).

setback for the forthcoming elections (2007b). The elections confirmed the prime minister, the Reform party finished ahead (Alas 2007a) of the social-liberal Centre Party and its candidate, who preferred a less controversial approach towards the monument. In March, Ansip's new government immediately laid the legal ground for the removal of the Bronze Soldier.

On April 26, Estonian authorities fenced off the central Tallinn statue. A day later, they removed the statue, exhumed the bodies of the Red Army soldiers underneath, and transferred them to a military cemetery in the outskirts of Tallinn ("NATO Sees," 2007). Unsurprisingly, the removal angered Russians, Estonia's ethnic minority and citizens of the Russian Federation alike. On the Russian side, the chorus of outrage was spearheaded by President Putin, who fiercely criticized the Estonian decision. In Tallinn, streets were filled with protesters against the decision of the Estonian government. Estonian police forces arrested hundreds of protesters (Adomaitis 2007). In the late evening of the day of the monument's removal, on Friday, April 27,[38] first signs of a cyberattack appeared on the monitoring screens of Estonian IT operators.

4.1.2 The attacks

Starting at around 10 p.m., Estonian organizations faced attacks on their servers responsible for handling e-mail, web, domain name resolution and other Internet services. Systems slackened or stalled under unusually high data traffic, Internet sites suffered from web defacement, email inboxes were filled with even more spam and phishing emails.[39]

[38] In their joint presentation, Gadi Evron, a known ICT security expert who arrived in Tallinn after the attacks had peaked, and Hillar Aarelaid, head of the Estonian CERT, speak of "Saturday, the 26th of April, 22:00" as the day, when the attacks started. A line later, they mention "Saturday, the 27th of April, 02:00" (Evron & Aarelaid, 2008). However, in 2007, the last Saturday in April was the 28th. In a post-mortem journal article, Evron (2008b) states that the attacks started at "10:00 p.m. on 26 April 2007". Street demonstrations that later led to riots took place on April 26 and April 27. Presentation slides made by Merike Kaeo, a U.S.-based Estonian security expert, contain a graphic of web traffic between Friday, 0:15 a.m., and Saturday noon, according to which traffic first abnormally increased on Friday night around 10:15 p.m., but culminated no earlier than on late Sunday, April 28. Interviewees confirmed that attacks started on a Friday, i.e., on April 27.

[39] For prior descriptions of the Estonian incident see also Evron 2008b; Herzog 2011; Landler & Markoff, 2007; Tikk, Kaska, & Vihul, 2010. Besides the sources mentioned in the text, findings in this article are based on interviews of the author with persons involved in the response activities back in Estonia and elsewhere in 2007.

Political institutions were an early target of the attacks. Estonian Prime Minister Andrus Ansip and other leading politicians were spammed (Berendson 2007). The email services of the Estonian parliament had to be temporarily shut down, as they were no longer able to handle the unusual data payload (Finn 2007). The Estonian news outlet Postimees Online fell victim to two DDoS attacks on its servers and had to close foreign access to its networks, thereby limiting chances of Estonian establishment to make their voices heard abroad ("Hansapanka Tabas Küberrünne," 2007). In addition, discussion forums of Postimees Online were spammed by bots with comments badmouthing and insulting the prime minister, making Postimees Online president liken the cyberattacks to an “attack on neutral and independent journalism” (Berendson 2007).

While defacements of governmental web sites may constitute an embarrassment for the sites' owners and symbolically undermine sovereign political institutions, they hardly constitute a major blow to a society and its security. The main causes for concern were the DDoS attacks on the Estonian infrastructure as they endangered the availability and functionality of services crucial for the Estonian society.

Internet traffic exceeded average-day peak loads by factor 10, (Aarelaid 2007) resulting in malfunction or non-availability of Internet services. Among the institutions affected was, most notably, the Estonian government. Its website, valitsus.ee, was not available for eight consecutive hours in the afternoon of April 28. For the following two days, response times often went up to a hefty 8 seconds and more, if the site was available at all. Statistics of Netcraft.com, a web site gathering information about up- and down-times of webpages, revealed that the website of the Estonian government failed to respond in 84 of 166 cases until Monday early morning (Hyppönen 2007c). Among the affected websites were those of the Prime Minister (peaminister.ee), the Ministry of Economic Affairs and Communication (mkm.ee), Ministry of Internal Affairs (sisemin.gov.ee), Ministry of Foreign Affairs (vm.ee), Estonian Parliament (riigikogu.ee), and, as already mentioned, the Estonian Government (valitsus.ee).[40] On Russian-language web forums, descrip-

[40] Hyppönen 2007b. Further domains attacked were: the Estonian Patent Office (epa.ee), Estonian Defense Forces (mil.ee), Estonian Academy of Music and Theatre (ema.edu.ee), Tallinn University (ehi.ee, tpu.ee), Estonian Business School (ebs.ee), Tallinn University of Technology (est.ttu.ee), a Yellow pages website (infoatlas.ee), a URL shortening service (zzz.ee) (Aarelaid 2007, confirmed in an interview with the author). Berendson mentions the following additional targets: “the University of Tartu, the Estonian Radio, Estonian Shipping Company, Woodman Pärnu furniture company, a real estate company Rime” (Berendson 2007). However, we have no statistically sound information about the effects on the availability of those websites.

Footnote continued on the next page.

tions on how to harm Estonian servers and Windows command shell scripts were published, along with pleas to run those scripts at a certain point of time.[41] Thousands of people running these scripts simultaneously results in web-traffic that over-stretches the capacity of those servers. This brief initial attack phase, which relied on humans executing the scripts, only lasted for a few days.

```
ping -n 5000 -l 1000 www.riik.ee -t″

@echo off
SET PING_COUNT=50
SET PING_TOMEOUT=1000
:PING
echo Pinguem estonskie servera, chtoby Internet u nih zavis!
ping -w %PING_TOMEOUT% -l 1000 -n %PING_COUNT%
sunic.sunet.se
snip out long list of targets
ping -w %PING_TOMEOUT% -l 1000 -n %PING_COUNT% ns.gov.ee ping
-w %PING_TOMEOUT% -l 1000 -n %PING_COUNT% mail.gov.ee
GOTO PING
```

Source: (2007, p. 12)

Figure 4.1: An example of a script used in the early phase of the attacks

In the second and main attack phase, the coordination of the attacks no longer depended on forum communication and synchronized human actions, but was mostly delegated to the command-and-control servers of real botnets. This phase started on April 30, lasted until May 18 and ran in four waves of different intensities, with different targets and attack techniques used. The “first wave” on May 4

Websites marked as available in Hyppönen's brief analysis were: Party of the Prime Minister (reform.ee), Ministry of Agriculture (agri.ee), Ministry of Culture (kul.ee), Ministry of Defense (mod.gov.ee), Ministry of Finance (fin.ee), Ministry of Justice (just.ee), Ministry of Social Affairs (sm.ee), Ministry of the Environment (envir.ee), Estonian Police (pol.ee).

Hyppönen’s analysis is ambiguous as to whether the websites marked as reachable had been attacked not at all, before or after the period of time analysed, i.e., for Saturday, April 28, 2007. In general, there is no consistent, conclusive assessment of the exact downtimes of the Estonian infrastructure during the entire three weeks of the attacks.

It is noteworthy that an attack on the web-services of an organization usually wouldn't affect its functionality. The attacks had “no impact on the Estonian military forces or national security apparatus” as a report by the U.S. based National Defense University holds (Miller & Kuehl, 2009, p. 3).

[41] Compare Aarelaid (2007). An example posted in a Russian website:
http://theologian.msk.ru/thread/list00350.php (last accessed in August 2012).

included DDoS attacks on websites and DNS systems.[42] The "second wave" on May 9-11 included DDoS attacks against mostly government websites and financial services. The "third wave" on May 15 included botnet-based DDoS attacks against government websites and financial industry. The "fourth wave" again consisted of attacks against governmental websites and banks.

Among the most significant attacks during this second phase were the attacks on Hansabank. Estonia's largest bank, recently rebranded to its parent company's name, Swedbank, owns a 50% share of national retail banking, which is almost entirely Internet-based in web-savvy Estonia. Its lending volume in 2007 was close to 7.5 billion EUR, its net profit in 2007 225 million EUR (Hansabank Group 2008). The web-interfaces for Internet-based services of the two biggest banks in Estonia were offline for about 45-90 minutes (Ottis 2009, confirmed by Estonian interviewees). The downtime period and limited availability amounted to losses of about 1 million USD (Landler & Markoff, 2007). On May 10, a day after the attacks on Estonian systems had reached their highest intensity, Estonian news outlet Postimees reported that Hansabank was offline that morning, that customers would encounter problems throughout the day, and that customers from outside Estonia would be denied access to the webpage.[43] Other than the attacks in the first phase, the second phase relied on botnets, which are regarded as the main vehicle and platform for cybercrime these days. The construction and use of botnets is based on divisions of labour. Botnets are created by so-called "bot herders" often using malware kits created and sold by highly gifted programmers. "Bot herders" then either sell their botnets or rent them out for a certain span of time to other parties, who can then use the botnets to send out spam e-mails, distribute malware, or, like in the Estonian case, launch DDoS attacks. The renting hours became visible by sharp rises of DDoS traffic at the beginning and like-wise steep falls at the end of a single attack (Frankfurter Allgemeine Zeitung 2007; Kaeo 2007).

4.1.3 Technical perspective of the attacks

As noted before, the cyber attacks on Estonia did not resemble a single ongoing steady campaign, but consisted of a number of distinct attacks over the course of almost four weeks. In what constitutes one of the more detailed texts about actual

[42] On these "waves", compare also (Tikk et al., 2010; Grant 2007).

[43] "Hansapanka Tabas Küberrünne," 2007. I used Google Translate to get a more or less precise insight into the content.

attack data and patterns, José Nazario, then researcher at Arbor Networks, a vendor for Internet security solutions, blogged about dates, lengths, destinations, and used bandwidths of the attacks. Between May 3 and May 17, 128 unique DDoS attacks on Estonian websites were counted, of which "115 were ICMP floods, 4… TCP SYN floods, and 9… generic traffic floods" (Nazario 2007). The attacks were unevenly distributed, with a mere three websites — Ministry of Finance, the Police and Border Guard, and co-hosted websites of the Estonian government and the prime minister — being the target of 106 out of those 128 attacks (2007). As to the bandwidth used, 94 remained below 30 Megabit per second (Mbps), 22 were located in the range between 30 to 70 Mbps and 12 between 70 to 95 Mbps. Regarding the duration of distinct attacks, 31 of the attacks lasted more than one hour, of these 8 lasted 5 to 9 hours and 7 lasted more than ten hours. Most telling however regarding the effectiveness of the attacks is that "10 attacks measured at 90 Mbps, lasting upwards of 10 hours" (2007).[44]

The discussion on the Estonian cyberattacks might make one believe otherwise, but from a technical perspective, the thrust and sophistication of the attacks were relatively modest, if not low compared to global standards even back in 2007. A survey of ISPs in the US, Europe and Asia on DDoS attacks conducted by anti-DDoS solution vendor Arbor Networks found: "In 2007, the largest observed sustained attack was 24 Gbps, compared to 17 Gbps in 2006. Thirty-six percent of the surveyed ISPs reported that they had observed attacks of over 1 Gbps in 2007." (Arbor Networks 2009, p. 2) In comparison, the Estonian attacks were modest (Cp. Clover 2009). Some interviewees from affected organisations even described the attacks and their effects on ICT systems as "boring". Given the overall throughput and capacity of the Estonian Internet, which was designed for a 1.4 million population, these attacks were nevertheless enough to obstruct the Estonian Internet infrastructure.[45] In addition, the attacks lasted far longer than typical

[44] Some Estonian experts doubt these figures, arguing that Arbor Networks, which was the supplier of at least one larger Estonian ISP at that time (US Embassy Tallinn 2007a), could only see a fragment of the Estonian Internet and therefore underestimated the amount of malicious traffic. There seems to be a minor fault with the figures in Nazario's blog. He first claims that there were only 7 attacks lasting 10 hours and more, and then "10 attacks … lasting upwards of 10 hours". In the table describing the durations of attacks, the numbers of attacks add up to only 126 attacks, leaving two attacks missing. Looking at the time intervals, it appears that this table misses the attacks lasting between 9 and 10 hours.

[45] A presentation by Merike Kaeo, doubleshotsecurity,com, provides some details on the topology of the Estonian Internet and government network (Kaeo 2007). The Estonian attacks showed that the low number and low capacity of international connections contributed to render Estonia's system unavailable. The connection of Georgian networks was even worse architected and therefore less resilient to cyberattacks, as the attacks in 2008 would prove.

DDoS attacks, not just hours and days, but weeks, albeit interspersed with periods of no or little malicious traffic (Duffy Marsan 2007).

Despite the lengthy duration, hiring a botnet for generating such malicious traffic would have been cheap. According to advertisements on Russian web forums, the costs for hiring a botnet for DDoS services for 24 hours and a bandwidth of 100Mbps was $75, the price for a week of 1000Mbps attacks was $600.[46] However, some security professionals involved in the response activities maintain that the attacks were technically and tactically more sophisticated and required a larger group of knowledgeable persons.

4.1.4 Costs of the attacks

The influx of DDoS packets had consequences on the quality and availability of Estonian web services — mainly the sporadic loss of services for government, communication, and banking (Ashmore 2009a, p. 8). E-mail and web services of some Estonian organizations were partly unavailable or functioned only at a reduced level. Government officials and journalists had difficulties in getting access to their emails or the Internet ("Estonia Hit by Moscow Cyber War," 2007). As one would expect for non-physical attacks like DDoS, the information technology structure was left undamaged, even though a "leading Estonian information technology expert" claimed that the attacks "'were clearly aimed at destroying the Baltic country's Internet backbone'."[47] According to security professional and researcher José Nazario, there have been "no apparent attempts to target national critical infrastructure other than Internet resources, and no extortion demands were made" ("Estonian DDoS - a Final Analysis," 2007).

Despite the press coverage and the political attention the attacks have stirred, a comprehensive post-mortem with a listing of precise downtimes and times of reduced service, aggregated and grouped per organization, complemented by a rough calculation of estimated financial consequences has yet to be written.[48] The lack of

[46] Segura & Lahuerta, 2010, p. 114. A previous version of the article with identical figures was presented at a conference in 2009, screenshots in the article capture advertisements published in September 2008. It is therefore safe to assume that prices for DDoS services were not significantly higher at the time of the attacks.

[47] Arnold 2009. According to a person from Estonia's cyber policy circles, the attackers managed to physically destroy a network component at an Estonian ISP.

[48] During the review process of this chapter, sources close to the Estonian MoD informed me that the Estonian Ministry of Defence had indeed conducted such a report, which would be about to be declassified. Alas not in time to be a source for this chapter.

data can be traced to the absence of overall monitoring of Estonian Internet systems in 2007 and the omission of systematic reporting by technical staff during the crisis. While the Estonian technical community still has abundance of data and log files, which could provide these answers (Estonian language skill provided), Estonian practitioners and international researchers alike obviously deemed such a study to be unimportant. Existing anecdotic evidence of damages occurred during the Estonian cyber attacks supports the conclusion that, despite shrill rhetoric heard in course of the events and in the aftermath, the financial losses more likely were "minimal" (Ashmore 2009a, p. 8). According to Rain Ottis, "only a few critical on-line services (like banks) were affected for clients inside Estonia", while "non-critical services (public government websites and news sites, for example) did suffer longer service outage" (Ottis 2009). The costs of the response activities, however, haven't been mentioned anywhere in the existing literature, neither the expenses for new hardware to scale-up existing systems or harden the perimeters of corporate networks nor those for over-time work of operational staff. Furthermore, according to an interviewee close to Estonian government circles, some banks accumulated substantial opportunity costs created by lost revenues.[49] A company's executive described the impact of delegating ICT staff to incident response tasks on ongoing ICT projects and the necessity to replan and reorganise these projects as the most prominent cost factor. Nevertheless, none of these costs should add up to figures creating greater public concern.

It is arguable whether the same can be said on the medium- and long-term effects of the relocation of the Bronze Soldier monument. The Estonian GDP numbers show solid, yet already decreasing GDP growth during the quarter of the attack. The trend of decreasing growth rates started several quarters before the attacks and continued afterwards into the financial crisis, sending the Estonian GDP into a brutal 14% nosedive in 2009 after an already displeasing previous year with a –3.6% recession.[50] The Baltic Times reported that Estonia's Transit sector income took a sharp dent in 2007, decreasing by 40% year-on-year. Russia, depending on ice-free Baltic harbours, diverted her cargo business from Tallinn's port to Latvia and Lithuania. According to an Estonian report and a Financial Ministry official mentioned in the article, Russia's economic payback aggregated to reductions in the Estonian GDP between 1 and 3.5 per cent ("May 9 Protestors Call," 2008). On

[49] I have not interviewed risk managers or persons with similar roles in banks that could have backed up these claims.

[50] Cp. data provided by Statistics Estonia, http://www.stat.ee/news-releases-2007. Also: Wikipedia's "Economy of Estonia" article has the real GDP figures, most likely extracted from annual CIA Factbooks. http://en.wikipedia.org/wiki/EconomyofEstonia.

the positive end, Estonia profited from a number of intangible and political gains. The attacks and the respective response turned Estonia into a household brand for all matters cybersecurity, which likely helped to secure the hosting of the NATO Cooperative Cyber Defense Center of Excellence and EU Agency for large-scale IT systems (European Commission - DG Home Affairs 2012) Its vanguard status was only increased by its support for some international cybercrime cases. Politically, Estonia managed to secure an increased commitment by NATO and the European Union, thereby advancing in its strategic foreign policy goal of strengthening integration into Western institutions and thereby balancing the influence of neighbouring Russia.[51] These issues lead over to the international and geopolitical implications of the Estonian cyberattacks, which probably have been more decisive than their effects on Estonian ICT systems.

4.1.5 The politics of cyberattacks

Soon after it had become obvious that problems with the Estonian Internet were caused by malevolent DDoS attacks,[52] officials in Estonia started blaming Russian authorities for being behind it. Ene Ergma, speaker of the Estonian Parliament, likened the attacks to "a nuclear explosion", the cyber attacks were "the same thing" (Poulsen 2007). The Estonian Minister of Justice asserted that some of the data packets in the flood lead to IP addresses belonging to Moscow offices of the Kremlin (Rantanen 2007). Prime Minister Andrus Ansip blamed the Russian government directly ("Estonia Hit by Moscow Cyber War," 2007). In an interview with a German daily a good month after the attack, President Ilves used slightly more contained wording regarding the role of Russia. He avoided calling it warfare. But asked how to label such kind of attacks, he said that, referring to potential unavailability of emergency lines, the attacks also "touched questions of life and death". Furthermore, he referred to the fact that Russian computers were involved in the attacks and that Russian intelligence service FSB would be able to control the Russian Internet (Frankfurter Allgemeine Zeitung 2007). Ilves continued, that some European states would have gone too far with their appeasement approach toward Russia.[53] Media representatives shared the view of Estonian incumbents. The edi-

[51] On Estonia's foreign policy options and strategies: Danckworth 2007.

[52] Some websites of Estonian authorities in addition suffered from web defacements, but they resulted more in symbolic damage than in actual economic or political costs.

[53] NATO later revised its policy towards the Baltic States in 2009, after Germany dropped its resistance to include them into NATO's defence and contingency planning ("Germany Behind NATO Proposal," 2010).

tor of the Estonian Postimees newspaper and website, Merit Kopli, talks decisively about the responsibilities: “The cyber-attacks are from Russia. There is no question. It is political.” (Thomson 2007) In her keynote at the FIRST conference in 2011, Melissa Hathaway called the Estonian cyberattacks in 2007 a “means of warfare”.[54]

The assumption of an immediate involvement of Russian authorities had soon been expressed by Estonian official and subsequently by scholars (e.g., Blank 2008; Grant 2007) who studied the Estonian incident and interviewed Estonian officials in the months after the attacks. Other researchers have subsequently joined in. Bumgarner and Borg unabashedly blame “Russia”, but provide no details about the specific role of Russian authorities.[55] Healey states “the obvious truth: the attacks were supported or encouraged by the Russian government and that to make the attacks stop, Western decision-makers needed to engage Moscow.” (Healey 2012, p. 2) Ashmore's detailed account on Russia's role in the attacks concludes that an involvement of Russian authorities had not been proven, but the mere belief of Russian involvement continued to frame Russian-Estonian relations until today (Ashmore 2009b, p. 8). Evron’s stance (Evron 2009) is typical for a representative of the technical community, and shared by many of the operational staff involved in the technical analysis and mitigation. Evron is reserved about blaming the Russian government, given the lack of straight evidence and a smoking gun. In contrast to the rhetoric used by some politicians and cyberwarfare theorists, technical experts shy away from calling the incident 'cyber warfare.’

Historiographic knowledge about the Russian policy during the events still is ambiguous and meagre. Gauging the involvement of the Russian government both in the attacks and their termination is difficult given the lack of sound and first-hand sources, Russia’s governmental records of those months still have the Cyrillic version of NOFORN stamp or higher. Lacking indisputable facts, judgement on Russia’s role is therefore mainly based on the perception of Russian foreign policy strategies, weighing of indications of Russian involvement, and the epistemological threshold at which pieces of circumstantial evidence add up to a picture “beyond reasonable doubt”.

[54] Gorazd Božič, “Melissa Hathaway keynote at FIRST2011 - Civil unrest is more like it IMO.” Juni 13, 2011, http://twitter.com/#!/gbozic/status/80205098986385409. Božič has been and is the head of CERT Slovenia.

[55] Bumgarner & Borg, 2009. The full report of the US Cyber Consequences Unit, however, has not been released publicly.

The assumption of the involvement of the Russian government and its close relationship to the unidentified perpetrators has been based on a number of arguments;[56] among them, that Russian and Kremlin IP addresses were involved in the attacks (Ashmore 2009a); that Russian experts had previously exerted similar attacks using the same botnets (Grant 2007, p. 6); that online and offline protests were coordinated;[57] that the scale and sophistication of the attacks required a serious organization behind the attacks;[58] the involvement of the Kremlin-directed Nashi youth group (Evron 2009; Grant 2007, p. 6); the long-term planning for the attacks;[59] Russia's asymmetric strategy against its increasingly West-leaning neighbours (Blank 2008, p. 230); that Soviet and Lenin tactics applied (2008, p. 230); the non-cooperation of Russian law enforcement agencies.[60]

[56] Partly compiled based on Mützenich 2009, pp. 8-9.

[57] According to an interviewee from Estonia's non-technical security circles, some of the organisers of the offline riots had been paid for their services by Russian intelligence services. An IT executive stated that local Estonian-Russians had likely opposed the riots, as a negative impact on the Estonian economy would be against their personal interests. Nevertheless, the Russian minority was highly susceptible to join in public demonstrations, some of which were distributively organised by snowballing text messages akin to techniques that later became popular in Iran or during the Arab Spring.

[58] An argument for sophistication instanced by members of Estonian policy circles is that the attacks included the hacking of a "key network device" of Estonia's Internet infrastructure. They had required dedicated knowledge of the Estonian infrastructure and resembled a "power demonstration of what can be done". One Estonian security professional described it as "targeted, single-packet router-killing stuff, never seen before". Another one dryly stated that there still was the possibility of pure chance that a router and its replacement got broken in quick succession. In addition, some hardware components are known to be vulnerable to so called "Packets of Death". Another example of sophistication mentioned was a sample of a bot malware that foreign security experts and police forces managed to get hold of on behalf of the Estonian CERT. International malware experts however told me that bot malware involved in the attacks was "not far beyond what had already been detected in the wild" back in 2007.

[59] Interviewees from Estonian policy circles stated that first signs of the attacks appeared way before the attacks themselves, among them very brief, intense floods of data packages to measure the capacity of the Estonian ICT infrastructure. The time span appears to be interpreted as an indication for strategic long-term planning by Russian authorities and as a counter-argument to the thesis of spontaneous online-riots by Russian nationalist geeks.

[60] (Evron 2008b, p. 124; "Venemaa Keeldub Endiselt Koostööst Küberrünnakute Uurimisel," 2008) Estonian authorities handed over a list of Russian suspects deemed responsible for the cyberattacks (an interviewee from Estonian policy circles: "We knew all the names of the criminals, we knew the masters"), demanding their extradition based on the Estonian-Russian mutual extradition treaty. The request was rejected by Russian authorities. An IT staff of an Estonian company stated that they had identified the "botmasters" and those "who organized these attacks" and that this information was then passed to the police. Other than in many other cases of cybercrime, the names of the suspects have never been publicised, though. Estonian authorities preferred this to remain low-key.

While these arguments have some weight, they don't add up to evidence "beyond any doubt". As the narrative in this chapter has shown so far, the attacks did not have a serious immediate impact on the Estonian society. What's more decisive politically are the long-term implications of the viability and utility of such cyberattacks. The cyberattacks would have fit into Russia's overall foreign policy strategy towards its neighbouring countries. Partly because of substantial ethnic Russian diasporas, partly because of security or national interests, Russia seeks to exert influence on the former satellite states it annexed in and after WWII and that gained independence after 1991. Its foreign policy strategy has aimed at containing Western influence in its neighbouring countries and the advance of NATO facilities towards the Russian border (Mützenich 2009).

What could Russia have gained by the attacks? The actual consequences of the attacks have been rather mild a) because of the existence of an Estonian cybersecurity community and b) its ability to timely link up to cybersecurity communities in neighbouring European countries and around the world. If these communities hadn't been in place, things had turned out differently. Given the still predominant ignorance of the role of global technical communities for Internet security and incident response among Western cyber security pundits, it is safe to assume that the attackers had not been aware of the Estonia response capabilities.

Without the latter, domestic politics would have been shaken up in Estonia. Had the attacks been successful, public and economic life in Estonia would have come to a standstill for days. After some time, probably a day or two, the technical experts would have found out what to do, how and with whom to collaborate, how to mitigate the DDoS attacks to bring ICT systems back to life. Much of the blame might have been put on the Estonian incumbent with his irreconcilable monument policy. His allegedly more Russia-friendly opponent, one of whose electoral strongholds lied in the Russian minority and who favoured a more diplomatic approach to the war memorial problem, might have gained a more favourable image among the Estonian electorate. Presumably more important than such an immediate gain would have been the long term effects. A successful attack would have left the impression among Estonians that Russia is capable to encroach into Estonian ICT systems and politics once it is fundamentally challenged by its neighbours' policies. Such an impression can lead to self-limitations in policy options; Russia would have increased its influence on one of the "near foreign countries".

From a political perspective, the strongest argument for an at least remote involvement of Russian authorities is the overall Russian strategy to their neighbouring countries and the tactics applied to decrease their neighbours' collaboration with and leaning towards the West. However, none of these potential gains materialized during or after the attacks. In this regard, whether the Russian government actually played a role in the attacks is a subordinate question. The political lesson is that cyberattacks can potentially be used as an instrument to fiddle with your neighbour's domestic politics.

However, given the lack of evidence for direct involvement of Russian authorities, this DDoS attacks might just have been a means of public protest like other DDoS attacks before. Pääbo observed a "failed integration policy" in Estonia's post-1988 nation-building process. "The bigger part of the Russian-speaking minority was treated as immigrants" and "presented as an undesirable relic of the Soviet period" (Pääbo 2008, p. 16). Earlier course of Estonian governments to not set up public Russian-language TV channels for their Russian-speaking citizens left Russian TV channels as the only source of audio-visual information for Russian-Estonians (2008, p. 16). Prime Minister Ansip implemented his campaign-promise and had the symbolic, yet highly controversial monument be removed. By that time, the Russian-speaking minority in Estonia —just as in other former Soviet satellite states — had been sufficiently frustrated by its marginalisation to be susceptible to calls for rallies both off- and online. The arguably mid-level technical sophistication supports the plausibility of the argument that the attacks were initiated and conducted by Russian-speaking civil society hacktivists renting out cybercriminals' botnets. Early verbal escalation by Estonian politicians helped to internationalise an inner-Estonian quarrel between the Estonian government and the Russian-speaking minority. The deterioration of relations between Russian on the one side and Estonia and Western allies on the other create a political climate promotive for deepening the integration of Estonia in the Western hemisphere at the expense of Russian influence.

The weakness of the existing historiography of the political dimensions of the Estonian cyberattacks is that it still lacks hard evidence. Only a look into the archives will reveal the actual considerations of the actors involved in the crisis.

4.1.6 Conclusion

The attacks on the Estonian Internet infrastructure had only relatively mild direct impact on the Estonian society. Sure enough, Estonian organisations and their IT departments bore the costs of delegating their staff to incident response tasks, political institutions' cultural capital was diminished by web defacements and other forms of ridiculing. The long-term relevance of the Estonian cyberattacks in 2007 is not that it allegedly constituted the first instance of a cyberwar. It did not when one uses a serious, sober definition of that term. Nevertheless, the attacks constituted a watershed in the history of Internet security for two aspects.

First, the attacks made it plausible to a wider public that cyberattacks can be used as a tool in international or bilateral conflicts. This feature is irrespective of how one answers the question of who was behind the attacks — whether it was a loosely-coupled, ad-hoc group of feverish Russian nationalists with varying degrees of IT skills plus some knowledge of how the cybercrime underground economy works; or whether it was a team within the Russian FSB collaborating with befriended cybercriminals of the Russian underground economy, connected with unknown levels up the ladder in Russia's security bureaucracy and administration. Irrespective of the answer, the attacks fitted well into the overall Russian foreign policy strategy towards its neighbouring countries at that time, which was characterized by an increasingly hard-line stance of the Kremlin and the drive to increase their cultural, political and economic influence in countries neighbouring her Western borders.

The Estonian political response, in concert with its Western allies, was to deter Russia and other countries from future application of attacks against civil Internet infrastructure against a country. A mix of diverse policy approaches has been implemented. Government representatives have rushed to name-and-shame state-funded or state-tolerated attacks on civil ICT infrastructures, branding it as illegitimate international conduct. Media coverage of the events has portrayed Russia's more dominant foreign policy in the near countries, exposing close relationships between underground economy, intelligence services, and government circles. A long-term endeavour has been to shrink the "grey zone" of arguable only-just-legality of aggressive cyber-conduct. On the technical-operational side, increased alertness and preparedness for such attacks has been a goal of policy makers.

Estonia and its Western security allies have assured their mutual support in the event of future, large-scale attacks on their ICT infrastructures, thereby raising the risks and potential costs for an adversary tolerating or even utilising voluntary groups to attack foreign Internet infrastructures. As a result, Estonia has become more embedded than ever into Western security and policy institutions, while Russia's cultural and political influence on Estonia has been further reduced. Whatever Russia's foreign policy circles, if involved at all, had defined as strategic goals, the Estonian cyberattacks hardly advanced Russia's political cause.

The second aspect is less obvious, more hidden, but nonetheless highly relevant both for future Internet security incidents and for questions of democratic governance of communicational infrastructures: the relationship between networks and hierarchies, between operators and owners of communicational infrastructures and traditional security institutions.

The Estonian cyberattacks will go down in history as a rare case where a Minister of Defence stated that his country was in a "national security situation" — and yet the relevant contribution to straighten out the situation did not come from military staff, but from a community of technical experts, which operated under the norms of the "beer & sauna protocol" (Hillar Aarelaid, cp. p. 128), fancies conferences that start at 2 pm with a morning pint, and values effectiveness over procedure and protocol. The response to the Estonian attacks was a wild success of technical security communities' principles of lose governance, trust-based information-sharing and technology-facilitated ad-hoc collaboration. At the same time however, it marked the end of the community's autonomy from state interference and regulation. The relation between the domestic security community and the political sphere in those days is aptly symbolized by the location of briefings of high-level politicians by the security community: they frequently took place in the offices of the Estonian CERT.

Cultural and communicational conflicts between the technical community and the political sphere had already emerged during the attacks, when pieces of seemingly contradictory information from different sources of the community added up to an unclear picture at therefore distrusting and even reproachful political boards. As a thoughtful member of the technical community put it: "Governments and institutions simply do not know how to communicate with the community. They do not know how to do it. They are not used to it." Therefore, according to another member, "the biggest problem we face in these events is communication between hierarchies and networks." As a consequence, the community was formalised as a legal body (the Cyber Defence League), the informal core group of the response team now acts as a formalised technical advisory body to Estonia's national security council, the CERT's hosting organisation Riigi Infosüsteemide Arenduskeskus (RIA; National office for information systems) has been granted special executive rights for future national security situations.

In an ideal world, such institutionalisation of the technical security communities helps to achieve two aims: to increase democratic control of Internet security governance and to increase the capacities and abilities of the overall response organisation against hostile intruders. Time will tell whether these approaches will serve the Estonian and other societies as well as, if not better than, the self-organised response of technical security communities in 2007.

4.2 Conficker botnet: Wasted talent for doom

No Internet security incident has raised the question of the scalability of the networked approach more assertively than the so-called Conficker botnet. In late 2008, a malware exploited a critical vulnerability of the Microsoft Windows operating system, installed itself in a hidden section of the operating system, and propagated rapidly and silently to millions of other machines. Infected computers thereby became part of a botnet that progressively increased in scale. The botnet used an impressive variety of mostly known, but partly highly innovative features to propagate the botnet, to receive updates in a command-and-control structure, and to make the botnet resilient against rival criminals and defending security staff that would try to take it over. Despite its unusual size of roughly 12m bots at peak time, the botnet has only been used in a few minor cybercrime cases, making its underlying purpose mysterious to this day.

The Conficker case is emblematic of Internet security in several ways. The incident not only raised awareness of the problem of botnets among a wider global audience. It also featured an impressive global collaboration among technical experts and shed light on usually hidden security institutions that have kept the Internet running in times when the technical infrastructure has been attacked, and claims that the Internet is an anarchic place have been deeply woven into political rhetoric. In addition, the botnet gave a glimpse of current and future contradictions and problems for providing common infrastructural security and Internet-based aspects of national security.

The Conficker botnet has been covered in numerous publications not only at the time of its occurrence, but also in its aftermath until today. Thorough technical analysis has been shared in reports written by research institutions, the ICT industry and independent researchers (Porras, Saidi, & Yegneswaran, 2009/2009a, 2009b; Microsoft 2009a; Symantec 2009; Leder & Werner, 2009). Organisational issues are touched on in post-mortem reports conducted by ICANN staff, Microsoft or commissioned by the US Department of Homeland Security (Piscitello 2010; Microsoft 2009a; Rendon Group 2011). Apart from a current of brief, often semi-informed sensationalist news articles, journalistic accounts have substantially contributed to the Conficker narrative (Giles 2009; Markoff 2009a, 2009b; Bowden 2010, 2011).

This section gives a descriptive account of the rise of the botnet, the damages it inflicted, and the response to the attacks. It outlines the different activities and services that have been vital for the response to the botnet. This narrative is preceded by an analysis of the technical and institutional circumstances in the run-up of the attacks.

4.2.1 The rise of botnets

The Wintel era[61] has been plagued by a myriad of bugs[62] in its defining Windows operating systems. In the 1990s, software flaws only plagued individual users with the famous BSODs, the blue screen of death, named after the appearance of a computer screen that would inform the user about a software problem and the need to reboot their Windows machine. With the rise of networked computing and primarily the Internet, pervasive software bugs turned into mass vulnerabilities that could be exploited by nasty script kiddies, explorative hackers, criminals, criminal enterprises, and state intelligence or military agencies.

Windows' numerous vulnerabilities have spurred the rise of countless pieces of malware. Among the most prominent have been worms, named for their seamless, automatic propagation, which requires no manual intervention. The Morris worm, the first known instance of this breed of malware, infected between a twentieth or a tenth of all 60,000 systems that comprised the very early Internet (United States General Accounting Office 1989/2013, p. 113). The infection ratio was comparable when worms like ILoveYou (also called VBS/Loveletter), Blaster (Lovesan), or Sobig hit the scene. In the early 2000s however, such an infection ratio equalled millions of infected machines. With the exception of 2008, every year until Conficker came along Windows systems were hit by a worm leading to millions of infected machines (Microsoft 2009a, p. 15).

Another remarkable trend of Internet insecurity in the Noughties was the rise of professional criminal organizations behind hacking and the use of botnets to serve their ends. Online crime has "taken off as a serious industry since about 2004"

[61] It has arguably come to an end, after all. Mary Meeker of Silicon Valley's venture capital behemoth Kleiner Perkins Caufield Byers observes a demise of the market share of Windows systems among personal computing platforms from more than 90% from the early 1990s until 2006 to 35% in 2012, caused by the rise of Apple's iOS and Google's Android (Meeker 2012, p. 24).

[62] In the early pre-silicon days, computers consisted of tubes wired in space-consuming circuits. The air for cooling these constructions was at times not filtered adequately, so that eventually bugs would fly or crawl in, land on one of the one of the delicate, hot tubes and thereby extinguish not only their life, but also that of the tube. The computer would start to malfunction; operators would look for causes and eventually find the remains of a vaporized bug on or underneath a tube with dark debris on the glass inside the tube. That's why errors in computing devices, hardware failures or software errors, are called bugs (Shapiro 1987).

(Moore et al., 2009, p. 3). The cybercrime scene has turned into a networked underground economy, which is highly decentralised and primarily consists of individuals and smaller groups offering certain services necessary for conducting Internet-based fraud, theft, extortion, identity fraud or denial of service (Steigerwald et al., 2011; Laufen 2010; Fossi & others, 2008; Stone-Gross, Holz, Stringhini, & Vigna, 2011; Goncharov 2012). Furthermore, traditional underground mobs are deeply engaged in cybercrime and have therefore acquired sophisticated technical expertise (Menn 2010). According to a member of a German police high tech crime unit, this networked underground cybercrime economy consists of "at best some 1000 to 5000" criminals worldwide (Laufen 2010). Botnets play an important role in this cybercrime economy. They are the platform on which all sorts of cybercrime are launched or hosted, including sending out spam e-mail, hosting criminals' webpages used in phishing attacks, or launching DDoS attacks.[63]

The roots of botnets trace back to Internet Relay Chat (IRC) systems and their management. To go beyond the core chat features and to automate administration, operators and enthusiasts expanded IRC functionality by adding scripts (Canavan 2005, p. 4; Schiller et al., 2007, p. 79). The notion of "botnet" is apparently based on the 1993 "legitimate IRC bot" called *Eggdrop* and its so called "botnet" feature, which allowed IRC operators to "link many instances of the bot together" (Canavan 2005, p. 4). By the end of the decade, the first malicious IRC bots had appeared, which would allow botherders to control the infected machines (2005, pp. 6-7). These early botnets already employed a wide range of features used in subsequent botnet malware and covered propagation (luring users into installing malware, scanning target machines for exploitable software vulnerabilities, exploiting other malware, guessing or attacking weak passwords), resilience (anti-Anti-Virus, encryption), communication (talking with C&C server via IRC or other channels), and actual botnet services (ranging from DDoS, spam, phishing, storage and distribution of illegal or unlicensed content, data mining, ransomware, or adware installation) (Schiller et al., 2007, pp. 30-62). Around 2002, the increasing number of early IRC-based botnets was perceived as a "growing problem on the Internet" (Baylor & Brown, 2006, p. 3).

The centralised nature of IRC-based botnets made them vulnerable to decapitation by taking down C&C systems. IRC-based botnets rely on IRC servers to get the message from the botherder to the bots. Industry and law enforcement landed some early victories in fighting IRC-based botnets by fostering collaborations

[63] Moore et al., 2009, p. 5; cp. also the case of the Estonian cyberattacks in this chapter.

among ISPs, law enforcement, and the security community. Botherders and bot malware authors tried to evade this situation by adding various DNS techniques to their botnet architecture. Using domain names instead of hard-coded IP addresses, and special DNS-related hosting techniques such as multihoming, dynamic DNS and fast-flux DNS increase the botnets' resilience.[64] The latter proved to be a particularly effective technique to make "taking down botnet C&Cs … impractical to a large extent." (Schiller et al., 2007, p. 91) With fast-flux, the IP address of a domain is frequently and rapidly changed in the DNS system, making a botnet's C&C systems resilient against IP blacklisting. Hence malware authors and botherders shifted to fast flux DNS (2007, p. 24) — a move that left some in the security community in alarmist agony: "If left unabated, the botnet plague could threaten the future of the Internet, just as rampant crime and illegal drug use condemn the economic future of real neighborhoods." (2007, p. 25) With botnets hiding behind the technical and organisational givens of the DNS system, established techniques to take down botnets had become ineffective by 2007, leaving "security professionals who pioneered the fight against botnets … demoralized" (2007, p. 2). With the emergence of fast-flux, the domain name system emerged as the battleground for bot authors and herders and the response side.

4.2.2 A technical-organisational exploit

After the devastating experiences with worms in the early 2000s and the first appearances of botnets that used propagation techniques similar to worms, it was only a matter of time until a truly wormable botnet malware would be unleashed. The time had come when Microsoft announced the release of a critical extracurricular security update in a security bulletin on October 23, 2008 (Microsoft Corporation 2008). Security Bulletin MS08-067 stated that the update would resolve CVE-2008-4250, "a privately reported vulnerability" in Windows (National Institute of Standards and Technology 2008/2008; Microsoft Corporation 2008). A bug in the Microsoft Server Service allowed an attacker to take control of the attacked machine and remotely execute commands in Windows 2000, XP, and Serv-

[64] The Honeynet Project provides an apt definition of fast-flux: "The goal of fast-flux is for a fully qualified domain name (such as www.example.com) to have multiple (hundreds or even thousands) IP addresses assigned to it. These IP addresses are swapped in and out of flux with extreme frequency, using a combination of round-robin IP addresses and a very short Time-To-Live (TTL) for any given particular DNS Resource Record (RR)." The extended version of fast-flux: "Essentially the domain names and URLs for advertised content no longer resolve to the IP address of a specific server, but instead fluctuate amongst many front end redirectors or proxies, which then in turn forward content to another group of backend servers" (Riden 2008).

er 2000; Windows Vista and Server 2008 were affected, but to a lesser extent (Microsoft 2009a, p. 41; Symantec 2009, p. 6). Despite Microsoft's increased security efforts after the previous worm disasters, the bug that was first introduced with Windows 2000 continued to infest Microsoft's operating system until Windows Vista (Microsoft SDL Team 2008; Sotirov 2008). Only a few days after Microsoft's security update, the Gimmiv.A trojan was caught by honeypots of AV vendors and other security-interested people. The trojan used a specially crafted Remote Procedure Call (RPC) request to exploit the aforementioned vulnerability in the svchost.dll, the file that implements the Windows System Service (F-Secure 2008; McAfee 2008). The trojan was designed to spy on users and harvest personal information from their infected machines.

On November 20, roughly a month after Microsoft's implicit hint at Windows' glaringly open hole, a more serious piece of malware made use of the RPC bug to creep into PCs and take them over.[65] The disaggregation of the ICT security and anti-virus industry has led to different names for the malware. "We're all sorry for that", F-Secure's chief researcher and talking head Mikko Hyppönen remarked (Hyppönen 2009, p. 4). Symantec innocuously dubbed it *Downadup*, Kaspersky came up with *Kido*, but the term *Conficker*, coined by Microsoft's security teams, eventually stuck. The word is apparently a play with the letters of a scareware briefly distributed by Conficker called Traffic Converter and an honest assessment of the blitz-worm (Bowden 2010). *Ficker* is the German for *fucker*. The malware did exactly what the security community was expecting: It exploited CVE-2008-4250, so that infected machines tried to contact and infect other machines on port 445 where the Windows Server Service was listening for inbound requests (Symantec 2009, p. 7). One of the first known infections happened on the early evening of November 20, recorded on a honeynet — a system of honeypots — operated by SRI International, and initiated by a machine with an IP address that is assigned to a network in Buenos Aires (Bowden 2011, p. 1).

The malware and the botnet it created used a remarkable and over time growing number of partly known, partly highly innovative features to propagate the malware, receive updates, and make the botnet resilient against rival criminals and defending security teams. The feature list of the different Conficker versions is rather long and includes: propagation using the MS08-067 vulnerability, Geo-IP data to determine language of target system, smart network scanning, brute-forcing

[65] While later dates have also been mentioned, the first traces of Conficker A date back to November 20. UCSD's Internet sink recorded 3000 infection attempts per hour in the late afternoon of Nov 20. Some fifteen hours later, that rate had soared to 115k/h (Giles 2009).

of simple SMB passwords in NetBIOS attacks in LANs, UPnP router pass-through, exploitation of Windows' AutoPlay vulnerability, HTTP rendezvous, peer-to-peer updates, security product disablement, date-based self-deletion, downloading of other payload, cryptographic verification of new instructions and downloaded payloads, pseudo-random domain-name generation, patching of exploited Windows vulnerabilities, push-updates via named pipes, blacklisting of certain IP-ranges owned by AV companies or Microsoft, sophisticated hooking approaches, dual-layer code packing, code obfuscation and code analysis evasion.[66]

Many of these techniques have been used in earlier malware. Conficker "stood on the shoulders of two decades of research and development" (Bowden 2011, p. 86). The scale, scope, and sophistication of its HTTP rendezvous technique and later the P2P techniques were new. While Conficker's propagation and infection mechanism relied on a serious Windows bug, the communication between bots and botherders relied on an exploitation of the organisation and governance of the DNS system, rather than its underlying technology. As a response to fast-flux botnets, the security community had started to cooperate with registries to get domain names blacklisted or blocked. The ball had been on the side of the botnet authors again. Then Conficker came. To avoid domain blocking, the malware and botnet designers utilized two technologies to make communication between bots and botherder more resilient: First, so-called HTTP rendezvous combined with a pseudo-random domain generation system and, second, a distributed system based on P2P communication. Bot-to-botherder communication based on HTTP rendezvous still relies on domain names and therefore a system that is not controlled by malware authors. The governance and technological control of the DNS system molds central and decentral elements. The DNS root is a centralized system; the DNS system with its TLDs is a decentral system; a TLD itself is centrally managed. To avoid law enforcement agency intervention, the botnet utilised at first modest, then large numbers of potential domain names. Thereby the attackers exploited the fact that until then no ad-hoc, speedy joint activities between all the actors in the DNS system had happened before. Cooperation between ICANN and TLDs only happened at the speed of quarterly ICANN meetings. Conficker C with its P2P techniques was even worse as it made the bot-C&C communication independent from central IP addresses or domain names. Communication between botherders and

[66] The technical description of the Conficker malware and botnet in this section and chapter is based primarily on various technical reports, such as Symantec 2009, pp. 51-52; Leder & Werner, 2009; Porras et al., 2009/2009a, 2009b, 2009/2009b, 2009a; Microsoft 2009b, 2009a, as well as news reports, blog entries, and expert interviews.

single bots happens indirectly over a chain of bots instead via one or several servers or domain names.

4.2.3 From A to C

Within a good month, Conficker A's "aggressive" (Symantec 2009, p. 9) propagation technique infected roughly a million PCs (Porras et al., 2009b), solely by exploiting the Windows Server Service vulnerability. This number is all the more stunning as Conficker A relied on the GeoIP service, a service which geographically locates a computer based on its IP address that was soon made inaccessible for Conficker A bots. The owners of the GeoIP service removed the data from the URL that was hardcoded into Conficker A source code, resulting in a decreased infection rate of Conficker A (Symantec 2009, pp. 21-22). This was one of the weaknesses of Conficker A to be removed in the subsequent B version. One of the decisive features of Conficker, which was then used by the response side to mitigate the propagation, was a technique called HTTP rendezvous. The basic principal of this technique is that the bots and the botherder meet at a certain address, here: a domain name, and at a certain time. Then and there, the bot, potentially receives orders and new malicious code from the botherder while the latter, possibly, gets some information from the former. To avoid easy detection by security researchers, date and time are not hard-coded into the malware, but generated by a domain generation algorithm. To make things as unpredictable as possible for the defenders, this algorithm is based on randomness. To let the bot and the command-and-control structure independently compute the same date and location, randomness is reduced to pseudo-randomness, i.e., the number of random outcomes is reduced by a variable known by the two of them, and only them. In the case of Conficker A, the current date was used as it is returned by major websites as part of a HTTP-adhering reply to an ordinary HTTP request (Leder & Werner, 2009, p. 11). The algorithm churns out 250 domain names per day dispersed over the five top-level domains .com, .net, .org, .info, .biz. The botnet was secured against hostile take-overs by security researchers or competing botherders by cryptographic verification of messages and payload signed and encrypted with keys only in the possession of the botherder. The chances of accidentally detecting Conficker by a user or even corporate IT operations were minimized by a variety of measures, among them the small footprint of the malware and the thrifty use of computing power and bandwidth. Despite all these features in place and the already speedy spreading, Conficker A was not used for anything — except luring a few users into buying scare-ware from trafficonverter.biz — but was replaced and supplemented by its offspring right after Christmas 2009.

On December 29, a month and eight days after Conficker A, Conficker B appeared on the scene, i.e., in the honeypots of researchers and AV vendors. Certain tweaks to the existing propagation mechanism, such as the removal of the restriction which meant it previously would not install itself on systems with Ukrainian keyboards, or the replacement of the web-based GeoIP location service with an inbuilt GeoIP table, led to an even further rise of Conficker infection rates. Propagation techniques like exploiting Windows faulty Autorun implementation or attacks on weak SMB passwords proved to be especially effective within corporate networks and the numbers of infected machines soared. The new version increased the resilience of the botnet against attempts to sever the bots from their command-and-control servers by increasing the number of TLDs with another three, namely .cn, .ws, and .cc. Starting on January 1, 2009, each Conficker bot tried to connect to an alternative daily set of pseudo-randomly generated 250 second-level domains dispersed of eight TLDs every three hours on any day. In addition, the bot malware authors introduced alternative paths for the botherder to communicate with infected machines. A new backdoor crafted into the Windows Server Service allowed the botherder to reinfect existing Conficker drones and thus to push new malware onto infected machines. Before replacing itself with the new arrival, a Conficker B bot would first check the validity of the incoming software package by verifying its encryption and signature. In an attempt to increase the botnets resilience, Conficker B blocked DNS lookup to sites of anti-virus and security solution providers, honeypots, and Microsoft. This measure prohibited any software on the infected machines from contacting these sites, e.g., to update Windows or download the latest virus signatures file. Unsurprisingly, Conficker B used a common anti-detection technique in malware, and disabled existing AV software and Windows' inbuilt security measures. Furthermore, to make life harder for reverse engineers, version B introduced anti-debugging and anti-reverse-engineering features. Strangely enough, Conficker B was apparently not used for anything. The multi-million-drone botnet giant was lying dormant. On February 20, the B++ version introduced a new P2P communication channel between infected bots via the Port 445 (Microsoft 2009a, p. 96). This port was the point of entry for the worm in the first place; on that port, the buggy Windows Server System component was listening for inbound communication.

The responding side did their utmost to isolate bots from its command-and-control rendezvous points (cf. the subsequent chapter 5, Producing Internet Security). Nevertheless, a few bots eventually managed to date their commander at the rendezvous domain and receive the instruction to update themselves with the new Conficker C version between early March and the middle of March (Porras et al., 2009/2009b; Rendon Group 2011, p. 22). Version C was a "significant revision" that left “as little as 15% of the original B code base untouched” (Porras et al., 2009/2009b). The two decisive new features, a significant update of the domain-name generation system and a new peer-to-peer communication unit, were a direct blow to the response efforts of the technical community. The new domain-generation algorithm increased the daily number of domains potentially used as HTTP rendezvous points to a whopping 50,000, distributed over 110 TLDs. Of these domains, a bot would randomly chose 500 per day and try to connect to them, leaving a bot with a 1% chance to successfully meet her commander (Leder & Werner, 2009, p. 9). This alteration of the DGA constituted nothing but a superb hack of the existing DNS system. Conficker C tried to achieve resilience for its HTTP rendezvous communication by exploiting the wide distribution of responsibility, authority and operational control in the overall DNS system and its governance. As an alternative route for bot-to-botherder communication, C featured a custom-designed, simple, but robust P2P protocol (Porras et al., 2009a). The P2P protocol allowed the botherder to push new updates onto the drones without having to rely on the domain name system that was controlled by the defending side. In addition, Conficker C further increased the resilience of bots by deleting Windows restore points on infected machines and improved its already sophisticated use of cryptology (cp. section “Implications of Variant C” in 2009/2009b). The infection rates of Conficker C were substantially lower than for Conficker B. Awkwardly, the malware authors had removed the propagation technique that used the CVE-2008-4250 vulnerability fixed by MS08-058 (Hyppönen 2009).[67] In addition, the response side had managed to block B bots from receiving updates from rendezvous points, only a few domains slipped through and allowed B bots to upgrade themselves to C. But again, Conficker remained a sleeping giant, hardly used for any malevolent activities other than creating and sustaining the botnet itself. In early April, a younger Conficker sibling, version E, briefly appeared on the scene and updated a small fraction of the botnet. It dipped its toes into the world of cybercrime by downloading other malware and scamware. It reintroduced the RPC attack that had turned the malware family into a worm in the first place. On May 3, the few E bots replaced themselves with the earlier, non-wormy C version.

[67] This is indirectly supported by a Microsoft report that analysed a data set consisting of telemetry data for a period of three weeks of telemetry data taken from a Conficker sinkhole. The report states that all version C bots, in MSFT lingo Worm:Win32/Conficker.D, used preexisting Conficker infections as propagation mechanism (Microsoft 2012, p. 7).

4.2.4 Impact

The worm turned out to be a smashing success in terms of its ability to reproduce itself onto millions of systems. On November 20, Conficker was first seen out in the wild. By the end of December, it had probably infected between 1 and 1.5 million systems with a unique IP address (Rendon Group 2011, p. 16). In its heyday, the malware was running on around seven million systems.[68] Conficker C has not been as prolific as its elder super-spreading siblings. At its peak in April 2009, more than a million systems with a unique IP address were Conficker C bots. As of early December 2009, this number had already gone down to between 300,000 and 400,000 unique IP addresses (2011, p. 10). In September 2013, Conficker C infections had gone down to about 11,000 unique IP addresses.[69] The Conficker family still continues to go strong and still is high in the charts of infections detected by Microsoft's security products. In last quarter of 2011, Conficker was detected on some 1.7 m systems (Microsoft 2012, p. 4). In 2011, Conficker was still the no.1 plague for corporate ID administrators. Leading with a 13.5% share and by a 5% margin over the trailing malware family, Conficker still topped the chart of malware found on infected organisational systems in 2011 (2012, p. 77). In April 2012, a Microsoft employee reported that Microsoft had removed "Conficker from approximately 283,000 computers per quarter on average for the past year." Nevertheless, the "worm continues to be persistent" (Blackbird 2012). In 2013, it was still at in third place out of the most common infections on corporate machines (Microsoft 2013, p. 76).

The actual figures are somewhat nebulous as infections are usually measured as unique IP addresses. Whether one unique IP address really represents one computer has not been thoroughly analysed in the literature available and somewhat disputed among scholars. Shadowserver has a warning on their website:

> "A single unique IP address may not represent a single infection. It could represent 100, or 1/100, and anything in-between, or even higher

[68] Microsoft 2012, p. 4. This is supported by Shadowserver Foundation's botnet analysts, cp. their website at https://www.shadowserver.org/wiki/pmwiki.php/Stats/Conficker. The newspaper Telegraph however reported, "Tom Gaffney of F-Secure estimated that 15 million machines were infected" (Adams 2009). As Shadowserver has had a crucial role in sinkhole operations and the analysis of the telemetry data, their word has more weight in this matter. Rendon Group choses the diplomatic path: "between five and thirteen million computers from approximately 6 million unique IP addresses are infected by the A or B variants of Conficker." (Rendon Group 2011, p. 10)

[69] Cp. CWG homepage, accessed February 2014. The CWG has apparently stopped counting the number of Conficker C infections in September 2013.

> or lower of the listed values. The purpose in giving any numbers is to have a starting point in treating all the values in an equal manner."[70]

The overall installed base of Windows was estimated to be around 1b in mid-2008 (Foley 2007).[71] The infection ratio for the Windows populace hence was, even at Conficker's peak, a mere 0.7 percent.

Botnets are the preferred tool for many types of criminal Internet-based activities. While Conficker has largely been lying dormant, it has been used a few times for malicious purposes. The few forays into real-world usage of Conficker were in the field of criminal activities. In its early days, users of Conficker A machines were sent to the infamous scareware distribution website trafficconverter.com. The later minor version E tried to lure users into downloading the scamware Spyware Protect 2009 and installed the Waledac malware (Krebs 2011, 2009; Gostev 2009; Microsoft 2009a, p. 105). In terms of criminal theft activities of which botnets are notorious, Conficker has been more about potential than actual damages. Luckily, the botnet has never lived up to its doomy promise.

The damages Conficker inflicted are hard to estimate. No figures about the direct costs of these forms of Internet-based theft have been published. But direct theft accounts only for one aspect of cybercrime related costs. A thorough assessment needs to consider all incident related costs of cybercrime including a) "direct losses and indirect costs such as weakened competitiveness", b) "costs in response to cybercrime" such as expenses for staff, software, hardware or service to mitigate an incident, and c) "indirect costs such as reputational damage…, loss of confidence in cyber transactions… and the growth of the underground economy." (Anderson et al., 2012, p. 4) As Conficker might as well have been designed for non-criminal, but equally shady purposes such as a military cyber-weapons platform or an intelligence tool, indirect costs would also include the long term societal costs of securitization, militarisation, or policing of the Internet. But these costs are unknown. No coherent study exists that neatly summarises the damages inflicted by Conficker or any other single botnet or malware. As is so often the case in the domain of Internet insecurities, knowledge about the costs is only cursory.

It is apparent however that Conficker has left its mark in some corporate IT departments, either inflicting over-time pay to internal staff or expenses for external support. Cleaning organisational networks has been a challenge as the still ongoing

[70] http://www.shadowserver.org/wiki/pmwiki.php/Stats/Conficker

[71] Annual shipment of new PCs have roughly been at 320m in 2013 (IDC 2013).

infections reveal. Some hospitals were reportedly affected by the worm in a serious way (Williams 2009). Infections within systems of traditional security institutions have also added national security implications. Several news reports about infected military networks popped up during the incident. Systems of the German army, French navy, British air force and navy were infected; reports of grounded French planes were later denied by French officials however ("Hunderte Bundeswehr-Rechner von Conficker Befallen," 2009; Bachfeld 2009; Willsher 2009). Computers had to be disconnected to hinder malware propagation and to block potentially illicit data transfer from the infected machines to an outbound location. Reports about sniffed passwords or sending out spam email appeared in respected ICT media outlets ("Hunderte Bundeswehr-Rechner von Conficker Befallen," 2009). Based on the findings of the binary analysis of Conficker, such behaviour of a Conficker bot seems implausible.

4.2.5 Reviews, framing

Just like the Estonian attacks in May 2007, Conficker became the threat du jour in early 2009 — a cyberinsecurity phenomenon that was to be mentioned in any cybersecurity discussion until it was replaced or at least supplemented by Stuxnet and the Wikileaks leaks and dramas. Despite their often well researched articles — geography, culture, and media tradition help when it comes to easy access to sources in ICT industry and security community — U.S. media reports applied the usual militaristic vocabulary, putting the sticker “cyberwar” on all matters of Internet security. John Markoff’s article categorized Conficker as “cyberwar” in the New York Times (Markoff 2009a), Mark Bowden’s article in The Atlantic opined that botherders and the defending side were engaged in a “cyberwar” (Bowden 2010). His well-researched book on the Conficker was even subtitled “The first digital World War” (2011). Two years after Estonia’s speaker of the parliament likened the Estonian attacks to a “nuclear explosion” (cf. section 4.1.5), German security scholar Sandro Gaycken mentioned “assessments” according to which the worm constituted the “first ‘nuclear test’ of cyberwarfare” (Gaycken 2009, p. 12). The press coverage on what would happen when Conficker C would start looking for instruction on rendezvous points on April 1, 2009, nurtured the impression of a dangerous emergency situation. Wired’s Ryan Singel criticised security companies for “hyping the latest threat, whether that be Conficker or the latest PDF flaw” (Singel 2010). President Obama took a more sober stance, mentioning Conficker’s wide propagation on millions of PCs in a speech on U.S. cybersecurity strategy, only to announce a continuation of a “relatively lax” approach towards the private sector (Greenberg 2009). Unsurprisingly, the anonymous person that reported the bug “MS08-067 Server Service NetpwPathCanonicalize() Stack Overflow (CVE-

2008-4250)" was awarded the "Pwnie for Most Overhyped Bug". The black-humorous accolade is awarded annually at the Blackhat US conference to "the person who discovered a bug resulting in the most hype on the Internets and in the mainstream media". And Conficker was dubbed as "InfoSec Press Full Employment Act of 2009".[72]

Those who spent weeks, if not months uncovering the internal procedures of the various Conficker versions, had nothing but praise for the abilities of the malware authors. The researchers at SRI, who provided some decisive reports for the response community, exclaimed: "Those responsible for this outbreak have demonstrated Internet-wide programming skills, advanced cryptographic skills, custom dual-layer code packing and code obfuscation skills, and in-depth knowledge of Windows internals and security products." (Porras et al., 2009a) Felix Leder and Tillmann Werner, then botnet researching PhD students at the University of Bonn lauded the worm authors' "use of the hooking approach [as] technically very impressive" and claimed it was the "first time… that exploit shellcodes tries to evade [libemu]", a tool frequently used by malware analysts (Leder & Werner, 2009, pp. 20-21). The Rendon Group reported that "several researchers described it as 'elegant'" (Rendon Group 2011, p. 5). Researchers at Symantec, whose analyses were apparently less in-depth that those of the two previously mentioned groups (cp. Symantec 2009), were less impressed by quality than by the quantity of the features implemented in the malware: "This was one thing that set the Downadup worm apart from its counterparts of the last few years — its technical versatility. These secondary tricks weren't new; there were just so many of them." (2009, p. 2) NewScientist author Jim Giles nevertheless called it "one of the most sophisticated pieces of malignant software ever seen" (Giles 2009).

The motivation of the botherders is still as unknown as their identities. Speculations about the botnet's hidden purpose encompass research accidents, a platform for cybercrime or cyberwarfare, training sessions, or a distraction for the security community from other botnets actually used for cybercrime. Discussions about the purpose of the botnet and the identity of its backers however lack evidence. Even a seemingly general statement such as Rendon Group's conclusion that Conficker was "not designed to promote a specific type of attack" (Rendon Group 2011, p. 13) is not based on hard facts. Six theories about the perpetrators and their intentions have been discussed. First, until the release of Conficker B, a theory was that Conficker accidentally escaped the labs of benevolent researchers or academics (Bowden 2010). The new version however proved that the Conficker authors were

[72] "Winners of Pwnie Awards 2009", http://pwnies.com/archive/2009/winners/

following the activities of the response community and reacting accordingly. Second, the usual designated use for a botnet would be as a platform for cybercrime. The few forays into real-world usage of Conficker were in this field. Users were sent to scamware websites, the malware Waledac was distributed via Conficker (cf. section 4.2.4). Third, the fear of many observers has been, and partly still is, that Conficker is designed as a platform to launch any kind of malware or even cyber weapons. The most obvious, but also less sophisticated example would be to perform DDoS attacks, e.g., against critical infrastructures or private companies to blackmail the infrastructure owners. Fourth, in 2011, the CTO of a research company with the daring name U.S. Cyber Consequences Unit, Jon Bumgarner, went public with the idea that Conficker was used to deliver Stuxnet, the computer virus which destroyed the centrifuges critical for enrichment of Uranium as part of Iran's nuclear weapons program (Finkle 2011).[73] The arguments against Bumgarner's theory appear to be more convincing, though (Parker 2011; Ferguson 2011). On the other hand, a well-respected interviewee backed the Conficker-Stuxnet link. Fifth, another theory is that Conficker was just a training session, a rehearsal to lift botnet authoring skills. Under this scenario, motivations would vary depending on the background of the authors, whether one assumes they were ordinary rent seeking criminals, military or intelligence cyber units. Sixth, if one assumes either of the latter, the botnet could also be an Internet infrastructure intelligence tool or a step in preparing the weaponisation of the Internet (Rendon Group 2011, pp. 13-14).

4.2.6 Conclusion

The Conficker botnet had substantial real-world impact, but different than one would have expected for a botnet of its size. Its usage for malevolent purposes is unknown. Theories have waxed and waned, but apart from some relatively minor cybercrime involvement, Conficker has been lying dormant ever since. No one knows for sure why and what happened to the botherder. Is it an abandoned giant or a sleeper waiting for future missions? It is left entirely to our imaginations as to what could have been done with it — or rather, what could be done with it. For Conficker is still around, as Shadowserver's infection statistics show.

The mitigating response was initiated by a group of West Coast Internet security professionals, first loosely supported by a network of trusted persons willing to

[73] Bumgarner's main argument is that Stuxnet attacked Windows systems using not only four zero-day exploits, but also a vulnerability previously also used by the Conficker worm.

share knowledge and data. With the rise of their workload and increasing breadth of the tasks, the group formalised its organisational structure and distributed responsibilities among the group members. A vital role was later played by more than a hundred TLD registries, which helped to secure the counter HTTP rendezvous strategy, i.e., the blocking of every domain.

The response effort, as the next chapter will show, reveals the remarkable flexibility of the "security community" in monitoring, detecting and analysing an emerging security risk. If one excludes remote bot-cleaning and brute-forcing keys for encrypting and signing payload on NSA-like computers as a viable alternatives, the Conficker cabal and its supporters have designed and mostly successfully implemented a response strategy that possibly helped to keep the botnet at bay. The response effort however also uncovered the limits defensive DNS: it does not constitute remediation or an end-game. It just contains the botnet or rather, the botherder. Technically, there are still ways for the botherder to engage the botnet or update it.

5 Producing Internet Security

"You've got a bunch of superhero vigilantes that are keeping the wolf from the door"
Steve Santorelli

In light of the destructive capacity of contemporary cyber attacks exemplified by the Estonian 2007 incident and the Conficker botnet, how has the Internet survived without a comprehensive technological and organisational security infrastructure? A complete answer will be provided in Chapter 6. In this chapter, the ground will be prepared for directly addressing this question by way of recounting how communities and teams of technical security experts did in fact respond in the two cases mentioned. These teams in responding to the attacks demonstrated considerable ingenuity, creating products and services that will also be of interest. The role of peer-production in Internet security can only be gauged by a review of the solutions that have been produced.

For each of the two cases there is a distinct section in this chapter. In these sections, a detailed narrative of the response activities is provided. If necessary, technical background information is amended to help to understand the response strategy chosen by the actors responding to the incidents. The analysis is in addition supplemented by a description of the organisational approach of the response activities. Each case is addressed separately in this chapter, so that a detailed account can be presented of all pertinent response activities. Some attempt has been made to avoid highly technical jargon to facilitate comprehension. All response activities will be set in the context of the overriding organizational approach.

5.1 Estonia

How did Estonia respond to what has been called "the first attack of its kind, directed against virtually the entire informational infrastructure of a NATO country"? (Clover 2009) This question has not been sufficiently studied. Our knowledge is largely based on brief presentations by a very small number of technical experts, who were either personally involved in the response activities or one remove from those actually involved (Kaeo 2007; Aarelaid 2007; Kash 2008; Randel 2008). Written from a legal perspective, the most inclusive study on the Estonian events so far by Tikk et al. remains rather tight-lipped on incident response on the operational level (2010). The account does not say much more than that the response "was coordinated by CERT-EE with the help of system administrators and experts both within and outside of the country" (2010, p. 24). The Evron analyses benefit somewhat from data gathered at ground zero but, unfortunately, at a time when the bulk of cyber attacks were over. (Evron & Aarelaid, 2008; Evron 2008b). More recently, Jaan Priisalu, now head of the Estonian RIA, has enriched the debate and shared his impressions of the events in a public discussion (Gomez 2012).

The following sections give a description of the actors involved in incident response, their actions, and their resources to re-establish the security of Estonian Internet services.

5.1.1 Emergence of the Estonian community

The attacks did not come as a surprise to the Estonian security community. Estonia's informal yet tightly knit community was on alert. "When there are riots in the streets, they will eventually go cyber", was an assessment shared by many in the Estonian security community. But they expected an attack only for late May 8 or May 9, Estonia's contested historical date. In the last week of April 2007, the head of CERT EE, Hillar Aarelaid, was even attending a cyber forensics course at the Templemore college of A*n Garda Síochána*, Ireland's national police. The Garda Computer Crime unit then led an EU-sponsored *Falcone* project to establish a European platform for cyber investigation training. Aarelaid recalls Friday, April 28:

> "I received a phone call on the last-evening beer party. The phone call was quite simple: 'It started.' I was a little bit surprised. 'So early? OK, I'll be there tomorrow.'" (Interview)

The expectation of some sort of cyber attack was not based on intuition alone. In mid-April the message was spreading within the Estonian and wider international Internet security communities that on Russian-language forums commenters were calling to low-intensity cyber-arms, trying to find comrades who would help to initiate DDoS attacks against organizational pillars of the Estonian society.[74]

Estonia had a well connected and prepared national ICT security community in place by the time the attacks commenced. As early as the late 1990s, banks began collaborating and exchanging information on cyberattacks. At first, Information System Security Officers (ISSOs) of at several banks cooperated, thereby violating regulations that outlawed the exchange of information between banks. Banks had a strong incentive for this kind of cooperation. They confronted a public that was not yet convinced about the usefulness and security of online banking (Interview 83). The ISSOs got a green light from their superiors to collaborate with their peers at other banks. Later on, executive levels tolerated and the parliament legalised this species of collaboration among banks' ICT security staff. It was obvious to the security experts that collaboration was absolutely necessary and that it would have to include banks as well as suppliers.

As such, it was crucial to recruit security experts from all manner of economic sectors. In the early 2000s, bank information security personnel teamed with peers from ISPs, telcos, energy and other major companies. By 2003, an informal group to protect the national critical information infrastructure had emerged. This group of ISSOs also lobbied for the creation of a national CERT to ease cooperation and information exchange (Interview 64).

> "We started to invite to our meetings also people from ISPs and telecom, and maybe a year or two later we also asked to join us people from energy companies and a couple of other, bigger companies. We started realizing that we had created a small working group, which basically in-

[74] Global network security communities learned about the call-to-arms in Russian-language fora before the attacks actually commenced, just like their Estonian peers (Gomez 2012). It took these communities some three weeks to establish direct communication channels. Reasons were unawareness of the other's existence, mutual lack of trust and issues that appeared to be more important than liaising to peer communities. Based on existing links to their Estonian peers, some European technical experts, however, shared their insights on ongoing scheming in Russian online fora with Estonian security staff already in mid-April, i.e., weeks before the latter were granted access to communication channels of global mailing-list-based communities. Apparently, some Russian web fora are constantly monitored by various Western actors interested in Russia-based cybercrime, malware, underground economies, espionage and other suspicious activities.

> volved critical infrastructure people. Without attention, we were starting to protect the Estonian national critical infrastructure." (Interview 64)

Before long, the Estonian informal ICT security group allied itself with the traditional security apparatus.

The development of a community of Estonian ICT security experts was facilitated by the introduction of the Estonian ID card with cryptographic functionalities built in and the introduction of Internet-based voting. With the advent of Internet-based voting and the i-Elections of 2004, a task force of security experts from ISPs, election authorities, police, intelligence services and others was put in place to prepare for cyber interference in the electoral process (Interview 87; Interview 64).

> "So we formed this team and the task was to secure [the] first Internet elections in the world." (Interview 87)

Hackers around the world delighted in identifying and exploiting the vulnerabilities of electronic voting systems and this had damaged the emerging cyber-voting business sector. But making matters even worse, a member of the Estonian Internet voting project went public that Estonia's voting system was only as secure as the PCs of its users (Sietmann 2005). The task force attempted to counter vulnerabilities via continuous monitoring in real time of the Internet. No major incidents have been reported for one of these elections to this date according to local experts. After its establishing, the Estonian CERT (CERT-EE) has acted as the lead organisation for election-related incidents and their prevention. CERT-EE was established in January 2006 as a department inside RIA, the national office for information systems.[75] Elections for the national parliament were held March 4, 2007, and again the election task force, which continued to exist next to and supplemented by CERT-EE, met. With the onset of DDoS attacks seven weeks later in late April, one observer noticed, "it was quite easy to get people together again" (Interview 87). The Estonians were prepared and had trained well. They understood the demands of communication between teams and how to divide labors (e.g. forum monitoring, supervision of firewalls, the oversight of suspicious web fora). But they had never faced such a broad attack.

[75] RIA reports to the Ministry of Commerce. In 2011, the authorities of RIA were extended and it was renamed to Riigi Infosüsteemide Amet (Interview 87).

5.1.2 Team formation and first responses

A good month after the national election was held without major technical security issues, the informal Estonian community was on alert, again. It thought May 8 the most likely date for a spill over of the offsite riots to the digital sphere: "We had everything ready." (Interview 87) Persons close to the Ministers of Defense and Interior were informed about possible DDoS attacks. Estonian intelligence was likewise informed — it is part of the Estonian Internet security information exchange system. Despite the lack of a central availability monitoring of national Internet services, it soon became obvious to technical operators in Estonia that the websites of a number of local institutions had fallen victim to DDoS attacks.

Four hours after the attacks had commenced, by 2 a.m. in the early morning of Friday, April 27, operational teams responsible for governmental servers had realized in mutual updates by telephone that some government websites were exposed to Internet traffic exceeding normal traffic by 100 to 1000 times. Servers could not cope with the enormous traffic, hence, operational teams decided to move websites to "well-defended" web servers scaled to handle the excessive traffic (Evron & Aarelaid, 2008). What had started as an operational IT security issue — after all, DDoS attacks are almost daily business — turned into a national security situation three hours later when the chief public relations person of the Estonian defense ministry stated around 1 a.m. on April 28, "We are under cyberattack." (Kash 2008) In the words of his superior, the Estonian Minster of Defence, Jaak Aaviksso, "It turned out to be a national security situation." (Cited by: Landler & Markoff, 2007)

This "security situation" was subsequently mitigated by the Estonian community of technical experts with, in the beginning, mild support from their international peers. When the attacks commenced, CERT-EE naturally evolved as the central hub for information exchange and coordinated some of the defensive measures of operational IT units in Estonian organizations. According to Lauri Almann, Estonia's then Permanent Undersecretary of Defence, "we put together a team of experts from our departments of commerce, communications, military and the intelligence community, led by Estonian CERT" (Kash 2008). The actual response community however was much larger and included Estonian and foreign actors, such as:

- members of the Estonian Internet security community from banks, telecommunication companies,
- the five largest ISPs that "actually matter" (Interview 87),

- the Estonian Educational and Research Network (EENet) that used to administer Estonia's TLD and the respective DNS system (Interview 87),
- the national crisis committee, comprising envoys from various Estonian ministries,
- Estonia's National Security Coordinator,
- traditional security institutions such as the Internet police units, intelligence, counterintelligence,
- the National Security Coordinator and representatives from the Government Communication Office, the Ministry of Foreign Affairs, and the Ministry of Defence,
- contacts at supply vendors,
- the international CERT community, primarily in Finland and Slovenia,
- and last but not least the AV and security industry (cp. Aarelaid 2007; Interview 39; Interview 87).

A key role in the technical response activities certainly had the Estonian CERT, the Estonian technical security community and the global Internet security community.[76]

Collaboration with domestic actors was eased not only by previous collaboration, but also by Estonia's unique situation. In a country with 1.4 million inhabitants and a good 400,000 of them gathered in the capital Tallinn, geographic proximity and naturally close social ties facilitate defensive ad-hoc collaboration. In the words of Evron and Aarelaid: “Anyone can get in a car and drag people out of bed.” (Evron & Aarelaid, 2008) The nordic sauna culture played a supportive role, too. While Nokia's then newly crowned CEO Stephen Elop had identified it as a factor that led to the demise of Nokia (Johnson 2011) before the mobile behemoth eventually tanked under his leadership, an institutionalised mix of booze, sweat, and nudity came to the rescue for Estonia. Meeting with their peers in hour-long gatherings with alternating sauna and beer sessions, dubbed as the “beer & sauna protocol,” has created a level of trust among Estonian experts that allowed them to collaborate seamlessly during the attacks.

On April 30, Estonian experts came together for a joint meeting, representing organizations such as ISPs, mobile telcos, operators of the Estonian TLD and DNS, banks, police, and envoys from the government's Security and Information

[76] Gadi Evron's (2008a) take on who were the decisive actors for responding the attacks: “The heroes of the story are the Estonian ISP and banking security professionals and the CERT (Hillar Aarelaid and Aivar Jaakson).”

Boards. (cp. also Kaeo 2007) This group met only twice in person during the incident, as most of the collaboration was done online via IRC, wikis, emails and private meetings.

The situation the Estonian response community faced was apparent: a massive DDoS attack against the systems of various Estonian organisations.

5.1.3 DDoS mitigation

Initially, the Estonian Internet security community took the incoming attacks as a series of distinct smaller attacks. Defacement of the Prime Minister's website ("he's a dick"), spamming of Parliament's web services, a DDoS attack on the Ministry of Defence — all these incidents were handled case-by-case (Interview 87). As isolated incidents, such attacks on organisational Internet systems were and are rather ordinary. Dealing with them is everyday business, the response to them not very challenging. But the number of such distinct incidents and their duration grew. There were too many such incidents at Estonian organisations. The goals of the response effort — ensuring that all systems were up and running, delivering their usual range of services, and keeping downtimes as limited as possible — were unachievable. Almost two days into the attacks, the response strategy was adapted to meet the means available:

> "First, do not handle it case-by-case anymore. Second, go to sleep." (Interview 87)[77]

The choice of the adequate response strategy is constrained by the technical nature of the attacks, and the technical and organisational capabilities of the response side. There are two fundamental types of Denial-of-Service attacks: overstretching the capacity of a system by sheer amount of data or requests or, in a more qualitative approach, by designing requests in a way that the attacked system needs to spend many CPU cycles on the task. Responding to each requires a different approach. The scale of an attack is another factor to be considered for the response approach. Upscaling systems is an intuitive response to quantitative DDoS attacks. Another approach is to reduce the quality of one's services, so that a service is still delivered to many, albeit in a reduced quality. Stripping down a website and replacing it with a minimalistic version is an example of such an approach. Then there is the filter-

[77] "Sunday evening 17:00, we made the decision that we will no longer count targets. There is no point." (Interview 87) It's not quite sure though which Sunday he refers to. I think it's the first Sunday, but could also be the second, i.e., May 5.

ing of traffic directed at the attacked systems. The filtering can happen at any point between attackers and the attacked. With an increase of the scale of an attack, it is necessary to move the point of filtering further away from the attacked systems. And the scale of the attacks on Estonian systems increased.

Estonian organisations implemented a range of proven anti-DDoS measures, for it was not the first large-scale DDoS attack they had had to confront. Larger organisations like banks have the equipment and knowledge to handle smaller DDoS attacks (Interview 64). All they need to do in such cases is to reconfigure their systems, e.g., by adding new filters. Some also had dedicated anti-DDoS devices installed at perimeters of their networks. Intrusion detection and prevention systems help to detect malicious packets designed to overstretch the capacities of their network components and web systems (Interview 64). Even when bank systems came to a halt, the response was dry:

> "We simply started to filter it… It was quite boring actually, they rebooted, they reconfigured the machines." (Interview 83)

But filtering on the organisational level was not sufficient to ensure the continuity of their IT operations. Entangled in a "kind of a war mode", one of the banks decided to focus only on their most valuable customers.

> "And the trick is, while we do this filtering, we don't have to communicate with everybody on the Internet. There is no point. You have to communicate just with the customer, with the people that you have something to do with. Most important customers have account managers. So you know where they are, you'll do your filters accordingly." (Interview 83)

Reconfiguration does not always do the job. New iron was also added to Estonian networks. Preparing for the expected onslaught on late May 8, a 10 Gbps cluster was set up to support ICT services for governmental systems. This effort, led and implemented by CERT EE, was supported by network equipment manufacturers that voluntarily lent their hardware in non-office hours despite the absence of contractual agreements or service level agreements that would have required such prompt support (Interview 87).

Eventually, the primary line of defence had to be moved away from the IT departments of affected organisations. Filtering inbound traffic at Estonia's upstream providers was a more viable path. The basic strategy in a DDoS attack is to stop malicious packets already at upstream providers.

> “[T]he ultimate goal with any security issues, especially large Denial-of-Service attacks is that you want to stop the traffic as close to the source as possible.” (Interview 69)

For such an approach, collaboration with first Estonian and then international network providers was inevitable. Upstream providers needed addresses of Estonian targets, analysed the traffic going through their networks to target addresses. When they see indications of malicious intent, they “try to limit that, …try to clean the traffic”, or establish a threshold for maximum traffic “just to give some breathing space” for the affected Estonian sites (Interview 65). Estonian ISPs tried to create such “breathing space” partly by blocking any traffic from abroad going to web services that were being attacked. While this nurtured the impression among foreign observers that an entire country was brought down, Estonian could still enjoy services that were mostly directed at them in the first place.

A more sophisticated technique than bulk-blocking of inbound foreign traffic is to find precise patterns by which attacking packets can be distinguished from normal packets. Network providers could then identify and drop such packets into whichever segment of the Internet between two communicating ends in which they occur. Distributing such attack descriptions up to the traffic chain is in the best economic interest of all providers, the networks of which are used to conveying large amounts of malicious traffic to the target address. Network expert Bill Woodcock elaborates:

> “But as your service provider, [the bad traffic is] still consuming resources in my network, it's coming to me and I throw it away. You are not paying any extra for that. So I take that attack description and send it to each of my peers [i.e., network providers with which the service provider has a peering agreement; AS] who are sending me that traffic because they are also not getting paid for the attack traffic, so it's in their best interest to filter that traffic. … It's usually not coming from a single source, but from a botnet. But you usually don't go all the way back to each source, at that point it's so diffuse that it doesn't make sense anymore.”

Identifying the difference between good and bad packets is a clear technological challenge.

> “If I'm your ISP, I just notice more packets. I can't tell whether you had a successful ad campaign or whether this is an attacker. … You need to define the difference between a normal packet and an attack packet.

> … If you can communicate that distinction to me, I can implement a rule in my router that will block the attack traffic so you no longer receive it.” (Bill Woodcock)

This is where a capable CERT can come to the rescue as they might be in a better position to identify these bad packets (Interview 23). Increasingly though, vendors of network devices provide solutions that can fulfil this task.

Heavy DDoS attacks are usually delivered by botnets. The Estonian attacks were no different. There were some attempts to understand the botnets involved in the attacks. As a strategy against a DDoS attack however, disrupting a botnet is not a feasible approach for an ongoing attack. Once a botnet has some degree of malicious sophistication, the time span it takes to understand the intricacies of the botnet malware, let alone to implement a proper mitigation strategy, is way too long to be a viable anti-DDoS approach. Furthermore, anti-botnet research was in its infancy in 2007. Continuous tracking of botnets was only started after the Estonian incidents by the likes of Shadowserver, Arbor Networks or Team Cymru (Interview 84). Consequentially, it is not known how many and which botnets were used for the attacks (Interview 84).

While this study focuses on the responses to incidents by communities of technical experts, the response also had a political dimension. The activities by politicians has had little, if any effect on the effective technical response by the Estonian and global Internet security community. The Estonian and Western policy sphere mainly responded with finger pointing at Russian authorities and a series of blame games and diplomatic arm wrestling. It was not and has never been proved that Russian authorities were involved in the attacks, and, if so, at which level. But the high-level political response signalled to those who might be considering Internet bullying that they would have to face political consequences. Whether this has had any impact on those who rented botnets for the DDoS attacks and engaged in some targeted attacks is only speculative. Not speculative however, but inevitable, was the support of actors outside of Estonia in managing the incoming torrent of malicious packets.

5.1.4 Communities' rapprochements

Once the attacks entered the second phase and increasingly became botnet-based, international collaboration and coordination became necessary.[78] The Estonian

[78] For an account on the different phases of the attacks cf. Tikk et al., 2010.

Internet security community had to collaborate with their global peers, first among them CERTs and operational network communities. The Estonians had not established relationships as close and trusting as necessary for seamless incident response collaboration with either of them. They all had to improvise.

Before proceeding with the description of events in Estonia in May 2007, a brief introduction to CERT communities is required. Computer Emergency Response Teams (CERTs) or Computer Security Incident Response Teams (CSIRTs) function as operational units to support individuals and organisations to adequately respond to a security attack.[79] Such CSIRTs are set up by large corporations or political bodies willing to consolidate responses to security incidents, to establish support centres for affected organisations or, in the case of corporate CSIRTs, organisational units. Different CSIRTs cover different organisations, sectors, regions, and countries. Their communicational cultures differ widely. While academic CSIRTs are accessible and communicative, staff of military CSIRTs have generally shied away from answering non-small-talk questions, let alone taking part in research interviews. Many CSIRTs are publicly funded organisations, or rather teams in publicly funded organisations.

CSIRTs themselves build communities via umbrella organisations. The Forum of Incident Response and Security Teams (FIRST) is probably the institution with the widest geographical reach among them. FIRST serves as a hub for more than 210 CSIRTs worldwide, hosting an annual conference in different locations around the globe that attract several hundreds, if not thousands, of participants every year.[80] Next to information exchange by presentations, networking is among the key goals of FIRST'S annual conference.

> "When you do the organization for an event like this, you specifically look for hotels that have a lot of chairs, where 2, 3, or 4 people can sit together and work on a problem." (Interview 43)

Security communities are not purely virtual communities, they also rely on personal meetings.

[79] CERT is a registered trademark of the CERT Coordination Center (CERT/CC) of Carnegie Mellon University. Generally though, both terms are used synonymously. CERT/CC was established in 1988.

[80] Compare description of meetings purposes by FIRST itself, http://conference.first.org/about/index.aspx

Already in the 1990s, there were plans to establish a dedicated European CERT, code-named EUROCERT. However, interests of European states proved to be too diverse to pool their response activities in one continent-wide body. Despite the loud fanfares from Brussels politicians, the recently established EU-CERT merely serves as the response team for the EU bureaucracy. After the EUROCERT plans ran aground, the Amsterdam-based Trans-European Research and Education Networking Association (TERENA) set up a CSIRT task force (TF-CSIRT).[81] Until today, TERENA and its TF-CSIRT serve as a mentor for the establishment of new CSIRTs and as a hub for CSIRTs to establish collaborative networks with their European colleagues. Its Trusted Introducer programme liaises new CSIRTs and their staff to the established CSIRTs, warrants the trustworthiness of a CSIRT with its accreditation programme. The frequent TF-CSIRT meetings create networking opportunities. In general, the purpose of TF-CSIRT has been to facilitate networking among European response teams that allow them to exchange ideas, information and issue requests for assistance in non-binding ways. Within this network of European CSIRTs, governmental CERTs gather in their very own "close-knit group" (Interview 82). Likewise, CERTs with similar constituencies tend to have their own communities (Interview 14).

By early May 2007, the Estonian security experts had not yet developed deep trust relationships with global security communities, which so frequently dealt with large-scale DDoS attacks on the Internet. This omission required a belated correction given the necessity for global collaboration, which translates to the exchange of potentially delicate data and the cutting off pf users in the domain of Internet security. The Estonians first created a bridge to their peers in European CERTs. The head of CERT EE flew to Prague to attend the 21st TF-CSIRT meeting in Prague on — lucky timing — May 3-4 to raise the attention of the European CERT community.[82] Aarelaid gave a brief presentation at the closed Trusted Introducer meeting on the first day about what was going on in Estonia and how the European CERT community could support the Estonians.[83] Trusted Introducer is a programme of TF-CSIRT to accredit CERTs and CSIRTs. CERT EE, which

[81] TERENA gave it the CSIRT name specifically because CERT has been trademarked (Interview 14).

[82] Cf. meeting schedule (TERENA, 21st TF-CSIRT Meeting, http://www.terena.org/activities/tf-csirt/meeting21) and the respective meeting minutes (Meiroșu 2007a). Hillar Aarelaid and Kauto Huopio gave a follow-up later that year at the TF-CSIRT September meeting (Meiroșu 2007b).

[83] On May 2nd, the second physical meeting in Estonia took place. Allegedly, Aarelaid could not attend the meeting as he was heading for Prague (Kaeo 2007, p. 16). Given the good flight connection between Tallinn and Prague, this sounds somewhat questionable.

was in operation only for a good year, had not yet run through all the steps required to get accredited. But given the circumstances, Aarelaid could attend the closed session and give an account of the incident, its effects and likely causes (Interview 14). Even if participants had no direct control over network devices, CERTs' staff usually have close connections to upstream providers or regular ISPs, both of which are helpful in an anti-DDoS response.

Lucky conference timing came to the rescue a second time. The cooperation between the international and the Estonian communities was significantly eased by a long-planned RIPE meeting in Tallinn on May 7.[84] It was during this RIPE meeting that members of the Estonian technical community were eventually introduced to members of the global Internet security community and gauged as trustworthy. However, for this to happen, many gallons of kerosene had to be burned. In January 2007, Aarelaid flew to Seattle to attend a closed meeting of "one of those security operational trust groups" hosted by Microsoft in January 2007 (Interview 69). One of the anti-botnet vanguards apparently paved the way for him to get to that meeting:

> "One of the conferences they [the Internet Security Operations community] ran was hosted by Microsoft. Hillar sent me a message that he wanted to come, an email. And for me, Estonia was this place in Eastern Europe where all this cybercrime was coming from. And I didn't know what to do with it. I thought it was a joke. Some cybercriminal trying to come to the conference. But I decided to take a risk and I brought Hillar there." (Interview 15)

At this meeting in Seattle, Aarelaid met a person that would happen to attend the RIPE meeting in Tallinn and eventually be decisive in closing the gap in the networks of security experts.

> "He actually called me and said that [he wanted to ask me for] a favour.... [whether I could] do introductions for him because he hadn't yet been formally vetted into an operational security group. ... [H]e was trying, he was in the midst of getting vetted, but he wasn't yet. ... And so I made the introductions...." (Interview 69)

[84] The *Réseaux IP Européens Network Coordination Centre* (RIPE NCC) is one of the five global Regional Internet Registries (RIRs) and provides "global Internet resources and related services (IPv4, IPv6 and AS Number resources) to members in the RIPE NCC service region" (http://www.ripe.net/lir-services/ncc). The region encompasses countries on the Eurasian landmass minus those east of Iran and Kazakhstan.

This "trusted intermediary" introduced Aarelaid to two experts who had key roles in the community of network operations, and long-standing international industry experience. With access to these persons, the Estonian community gained support for Estonia's important upstream providers and could move the packet blockade further away from their systems (Interview 69; Interview 23; Interview 65). The RIPE meeting on Monday, May 7, was both the perfect and the last opportunity to bridge this gap in the informal response networks in an informal environment. With this newly achieved status as members of the technical community, a few Estonians could now send requests or lists with attacking IP addresses to mailing-list-based security communities such as NSP-Sec.[85] Network security professionals around the world that are members of such a list would then help to stop malicious traffic from flowing from their networks towards Estonian systems.

5.1.5 May 8/9 day

While the attacks stretched over a period of more than three weeks, they were not conducted with a continuous thrust at the same level. After May 3rd, the thrust against Estonian banks waned, and DDoS attacks almost came to a standstill. The Estonian response community could have enjoyed a couple of calm days, were it not for a small number of very brief but intense attacks, interpreted as a means to measure the bandwidths of the overall Estonian infrastructure and its branching, and therefore indicating that more was coming. The attackers threw different amounts of data packets, say two, five, or ten Gbps, at Estonian servers to measure their capacities. Once the response side saw their systems melting under such unusual amounts of data, they tried to upscale their systems or applied other means to increase the capacity. On the following day, the attackers would re-measure the bandwidth. When they found Estonian systems with increased capacity, they would in turn increase their attack capacity, to which the response could again answer by increasing their bandwidth and server capacities.

> "Both sides probably knew on the second day that we are playing a game." (Interview 87)
>
> "It's like cardplay. I make a move, they make a move." (Interview 64)

Both sides could increase their respective capacities by either physically scaling up their systems or by logically reconfiguring previously throttled systems up to their

[85] Bill Woodcock, networking professional, co-founder and research director of Packet Clearing House, shared more details on the role of NSP-Sec in mitigating global DDoS attacks(Gomez 2012).

maximum capacity.[86] This job was mainly performed by Estonia's larger ISPs, jointly coordinated with CERT EE. Then, head Hillar Aarelaid recollects:

> "You're sitting together with sec guys of 5 big ISPs here and you say, hey, this is going on, and you need to respond accordingly and I know you have your hardware and please configure your systems accordingly." (Interview)

Apart from these measurement battles, those days in early May were rather quiet. The persons on the response side were "burning in their own oil", nervously awaiting what would happen next (Interview 87). Russian chat rooms were still being used by some of the attackers, and still monitored by some of the response side. The historic May 9 — late May 8 local time — still was the most likely day for another wave of attacks.

> "We thought that 99,999% they will start at 23:00 because of Moscow time." (Interview 87)

In a coincidence, one of Estonia's oldest ICT security companies, *cybernetica* (cyber.ee), threw its anniversary party in the Estonian Drama Theater. The venerable, prestigious building at the outskirts of the Tammsaare Park is just a brief 5-minutes walk away from the RIA building in Ravala Avenue, where RIA and CERT EE are located. Many from the Estonian Internet security community agreed to meet at that party, avoid drinking too many beers and then walk over to the CERT premises at around 10pm, an hour early before Victory Day, May 9th, would begin in Moscow.

> "You can see …some system admin guys coming in with ties with very funny mood from theatre…. They do not wear ties usually. Everyone was funny, joking, everybody knew that there will be a big event right now." (Interview 87)

Peers from the Estonian community and new foreign contact persons, who had came to Tallinn for the ongoing RIPE meeting, also joined the *cybernetica* party and then later the get-together at CERT EE.

While the Internet security community was expecting some kind of a digital attack, they had no idea what exactly they would be confronted with. Accordingly, when the attacks started, they were on high alert, but only for a short time.

[86] For an account of these bandwidth battles see also US Embassy Tallinn 2007b.

> "Everybody thought: big bang. But then it was all routine: Routing the traffic, sorting out the attackers, reporting it to the international communities." (Interview 87)

Nevertheless, the attacks were significant given the capacity of the Estonian infrastructure. But the attacks were sufficiently mitigated. By mid-May, the attacks had faded away.

5.1.6 Other security services

The production of Internet security in the Estonian case consisted of a successful mitigation of a mix of DDoS attacks, web defacements, and other malicious hacking attempts. This overall service can best be separated into five distinct services delivered by various Estonian and international actors. Situational awareness was established by the monitoring of technical and social networks. DDoS mitigation helped to minimize the impact of the main attack technique used against Estonian systems. These were the two main components of the Estonian response. Others included policing, malware analysis, and the somewhat vague category of information sharing and distribution. The main technical defence approach was presented earlier in section 5.1.3. This section discusses situational awareness, distribution of information, policing, and malware analysis.

Given the state of technology in spring 2007, responding to DDoS attacks on several organisations required data sharing among these organisations or their ICT providers. In Estonia, no central monitoring system existed that would have kept an eye on the health of Estonian web-based services. Individual organisations monitored their systems, security and AV service providers their clients' systems, but the countrywide state was not directly displayed on the monitors of some security professionals in Estonia. But there were workarounds. The IRC rooms used by the Estonian security community served as an improvised monitoring system that displayed problems with certain segments of the network. When two or more members from one organisation dropped out at the same time, then their networks likely had an issue (Interview 87). Most importantly however, security experts shared their knowledge with their peers in the Estonian community, achieving sufficient situational knowledge in the absence of central monitoring. As with DDoS mitigation, larger Estonian organisations and especially ISPs had their network monitoring systems in place (Interview 82). Their picture was complemented by the global community, where major ISPs with "global outreach have a very good visibility on DDoS activity." (Interview 82) When the attacks started and new events occurred during the almost-four-week period, peers in the response community were briefed rapidly. At some point, the response community agreed on cancelling their established reporting mechanism, which required any organisation to formally report an incident to the national CERT (Interview 87). The absence of written documents is certainly detrimental for historians and public managers who want to give an account of past events or review the performance of their Internet security institutions.

A second aspect of situational awareness refers to a knowledge of the activities of potential perpetrators. Russian underground fora are monitored by commercial security service providers, individual researchers, and presumably some governmental agencies as well. These fora are places where hackers, carders, botherders, malware authors and persons with other roles in the digital crime scene convene and collaborate (Interview 73).

> "Dozens, if not hundreds — every single underground forum is monitored by a big number of people data-mining." (Interview 73)

Information about potential backers of the Estonian attacks was shared within the Estonian community.

Furthermore, security teams and CERTs also need to stay on top of the news stream on the latest current security issues. They keep an eye on a substantial number of news sources and observe various sources that collect "badness information"[87] like Clean mx, Spam Cops, Spamhouse, T-Shield, Abuse Helper, and other trackers. A national CERT can normalize and sanitize such information and send it to their constituency.

Conveying lists of attacking systems from the attackee to networks closer to the attackers has been an important element in the Estonian response. Within Estonia, the communication channels of the local security communities were used. For trans-border communication, CERT communities or mailing lists like NSP-Sec helped to convey information. ISPs were also contacted directly or via their respective national CERT (Interview 82). National CERTs often act as intermediaries to convey or receive information on behalf of their constituency, especially when affected organisations do not have direct access to these communities via an employee-member. In the Estonian case, the Finnish CERT FI supported their Estonian counterpart by preparing lists for the international community. Estonian organisations send lists of suspicious IP addresses to CERT EE, which then "semi-

[87] I've heard this phrase mentioned by several interviewees.

automatically produced this lists of attackers" (Interview 87). This consolidated list was then sent to CERT FI, which took care of international distribution (Interview 87). Eventually, the CERT FI was attacked, too, and busy defending itself, so that the Estonian CERT then lost its most staunch supporter. Luckily, the attacks had already passed their nadir, so it "was not a big problem" (Interview 87). CERT EE took CERT FI's automation scripts and managed to accomplish the task of information distribution itself.

A fundamental aspect of security production in general is policing, an aspect that had been ignored by technical Internet security communities in very early years. In constitutional states, sanctioning actors for malevolent activities is transferred to the state's legal system. Beyond and besides, social systems usually have their own mechanisms to sanction activities that are unwanted by parts of the population. Policing and sanctioning on the Internet has been made somewhat difficult by the attribution problem. Proving that an unlawful activity has been committed by a certain person is difficult when a perpetrator applies sophisticated anonymisation and camouflage techniques. Estonian police arrested an Estonian schoolboy who committed some very basic manual attacks, but actually only played a marginal role. Nonetheless, he was the only one arrested for involvement in the Estonian incident. Requests for investigative assistance to Russian authorities bore no fruits.

> "We also found out who are the guys, who are the botmasters. We knew who was organising these attacks, and we gave the information to the police. And the police also made the request to Russia, a legal help request against some Russian personas, they were all nicknames" (Interview 83)

The Russian authorities did not respond.

> "My best reading of this is that Russia probably did not have full control over most of these attackers, or at least the organizers. But it was not in Russia's interest to actually solve this. So they were quite happy to say, 'Sorry, we can't cooperate with your investigation.'" (Interview 39)

Despite the early stages of anti-botnet research, Finnish malware researchers analysed the malware, identified its control server, and handed their findings over to CERT EE. Malware samples could be isolated from infected machines with the help of law enforcement that seized DDoS-packet emitting drones. The analysis of malware in sandboxed environments showed that different botnets were involved. Some malware came without a C&C mechanism; a worm was unleashed, akin to the *Allaple* worm that plagued the country only months earlier, simply to do harm and DDoS Estonian Systems (Interview 73). An Estonian man with a vendetta against his former employer, an Internet provider, had created a worm that would, independent from any C&C mechanism, attack the websites of an Estonian ISP and two insurance company websites (Hyppönen 2007a; Leyden 2010).

In this section on the response to the Estonian cyberattacks, this case of Internet security production has been portrayed as a set of input resources, activities, and outcomes, embedded in a specific technological, political, and social context. On an aggregated level, the response resulted in an outcome that can be called "re-established functionality of Internet-based services in Estonia". While the analytical framework developed in section 3.2.1 depicted incident response as a series of processes with distinct input resources and outcomes, the Estonian response activities is at times closer to a ball of wool with intertwined strings of activities, that the clear-cut diagram in section 2.4.2. Nevertheless, elements of standard incident response processes such as monitoring, analysis, forensics, sanctioning, and mitigation have all been present. The decisive component of the response certainly has been mitigation or, more precisely, DDoS mitigation. It comprised a distributed effort to filter and block "bad traffic", upscale systems and increase bandwidths. Other processes had a less prominent role in the response.

The activities relied on the sharing of information in the form of logs, hints, rumours, and analyses, which allowed the actors to achieve a shared level of situational awareness and informed decision-making. Distributed information was shared within and across communities to achieve shared, common awareness of the situation; likewise were mitigating actions performed distributively, partly in a coordinated, partly an autonomous manner. However, this leads us to the discussion of the distributiveness, openness, and socialness of the response. This debate is picked up in Chapter 6. Whether the production of Internet security in Estonia is an anomaly, or has practises in common with security production, is examined in the following section regarding the response to the Conficker menace.

5.2 Conficker

With the stunning propagation of Conficker and the race between the botherders and the responding actors with their constant upping-the-ante, the response side eventually comprised individual researchers, entrepreneurs and hobbyists, anti-virus or security software vendors, Internet service providers, domain name registrars, and ICANN. The sophistication and high-number of features of the Conficker worm required an unprecedentedly networked response to stop its contagion.

5.2.1 Boot-up of the response

Just as the Estonian Internet security community was expecting some kind of attack in April/May 2007, Internet security circles heard the worm coming weeks before it actually appeared in their log files. Whoever had an interest in malware was intrigued by the news regarding the MS08-067 bug on October 23. When Microsoft releases an extraordinary, not pre-scheduled software update, the reasons must likely be extraordinary, too. And if this had not been a broad enough hint, Microsoft's emerging counter-botnet executive asked the nearly 200 experts of the biannual International Botnet Task Force meeting in Arlington on the same day to stay vigilant and have an eye on the forthcoming issue (Bowden 2011, p. 42; Kaplan 2008).[88]

Malware researchers immediately started digging into the bug. One can only speculate what happened in access-restricted chat rooms, but even on public mailing lists like DailyDave the technical details of MS08-067 were discussed, Microsoft's software patch dissected it and compared it with previous versions and previous bugs. Only days after Microsoft's announcement, technical ideas were exchanged as to how this newly derived knowledge could be used to build an exploit.[89] Malware researchers were not the only community intrigued by MS08-068.

Operators of honeypots or Internet sinks — the more catchy names for passive Internet malicious traffic monitoring systems — observed a sharp increase in net-

[88] More on the International Botnet Task Force below in section 5.2.3.

[89] DailyDave is a security-oriented mailing list with a focus on security events, vulnerabilities, security research, and products of Immunity, the company of list-owner Dave Aitel. (https://lists.immunityinc.com/pipermail/dailydave/; older archives for November 2011 and earlier: https://lists.immunityinc.com/mailman/listinfo/dailydave). Examples of such discussions: Dailydave Digest, Vol 39, Issue 7, 24 Oct 2008, and Vol 39, Issue 6, 24 Oct 2008 (both available at https://www.mail-archive.com/dailydave@lists.immunitysec.com).

work scanning activities on November 21, the date Conficker A appeared. Such a surge would then happen again on December 29 when Conficker B said "hello world".[90] Sinks and honeynets are well suited to detect DDoS attacks, the spread of Internet worms and malicious software scans (Aben 2009).

Several security vendors, research institutions like SRI International (Bowden 2010) and projects of volunteer groups run their own honeynet systems. They too observed a rise in traffic, caught the new malware, and tried to understand its behaviour and design. The newly arrived binary was dissected, and important functions reverse-engineered by Christmas. Of particular importance was the domain generation algorithm, which was exposed by several persons working for research institutes or AV vendors (Giles 2009; Symantec 2009, p. 22).

With Conficker's contagiousness, slowing down the infection rates was a top priority in the beginning. To rescue came a particular feature of Conficker A. Its propagation module included "a single point of failure" that could be exploited by the response side (2009, p. 21). Before Conficker installed itself on a target system, it made sure that it did not target a Ukrainian system. To locate itself it compared its target's IP address with the likely location of this address using the Internet-based geo-location service Maxmind.com. When their systems suffered from the DDoS-like requests, Maxmind moved the file to a different URL than the one hard-coded into the worm, thereby slowing down the infection rate of Conficker (Hyppönen 2009; Symantec 2009, p. 21).[91]

Another immediate response taken by persons with an interest in anti-botnet activities was to block those addresses that Conficker had tried to contact to download new files, presumably new instructions or new code. It was this strategy that defined the activities of the Conficker Cabal and then the Conficker Working Group. In the beginning, the activities were rather ad-hoc.

[90] An Internet sink is a synonym for a network akin to the Network Telescope operated by the University of California in San Diego. The Network Telescope "is a globally routed /8 network (approximately 1/256th of all IPv4 Internet addresses) that carries almost no legitimate traffic because there are few provider-allocated IP addresses in this prefix" (Cooperative Association for Internet Data Analysis 2012). The research organisation SRI International runs a /16 network honeynet (Bowden 2010), probably one of the largest worldwide.

[91] Alas, the Symantec report, which actually is a collection of blog entries, does not explain how this actually slowed down the propagation of Conficker. The technical reports by Porras et al. (2009a) and Leder & Werner (2009) don't elaborate on this argument, either.

> "And then, based off what the threat was, [we were] trying to figure out how to take control. So we chose the DNS method. We will just file all the domain names. Using money…, money, throw money off us." (Interview 48)

At various ends, people picked up different tasks, voluntarily, not centrally managed, but somewhat coordinated using existing communicational channels, i.e., established security communities. How the strategy and the organisational approach came into being will be outlined in the two subsequent sections.

5.2.2 Response strategy

The response to botnets is defined by technology, norms, institutions, resources and time available for the responding actors. The response to the Conficker botnet was begun by an initially small network of security professionals, who chose to leverage their influence on elements of the DNS system to mitigate the attackers' exploit of the particularities of the existing DNS governance regime.

"Microsoft's initial goal was to slow the infection rate down to 'give time' to Windows users to patch their computers or utilize anti-virus software." (Rendon Group 2011, p. 16) That strategy apparently fell apart when societally critical systems were infected big time, infection rates soared, and the authors released Conficker B (2011, p. 16). The new version signalled that the attackers were closely following the response side's activities.

Operational control over one the key resources used by Conficker, domain names, was widely, globally distributed. Since the late 1990s, the DNS governance regime is characterised by a mix of technical and political provisions. Resolution of Internet domain names is based on the Domain Name System (DNS), a "cluster of functions" (Mueller 2002, p. 6) composed of technical protocols, Internet based services and governance institutions. By design of its protocols, the DNS is a federated hierarchical system. In principle, the DNS system is fairly simple. At the top, or rather the root, of the systems, there is a "root zone server" with a "root file" that contains a list of all generic and country-code top-level domains and the addresses of their respective DNS servers. For each of these top-level domains, there is another DNS server, containing a list of second-level domains under the respective top-level domain. The root zone is governed by ICANN, the data for the DNS root zone file is administered by ICANN's IANA department, and technical operational implementation lies with several root zone server operators. Individual ccTLD DNS systems are usually governed and operated by national registries, while generic TLD's are governed by Verisign, Neustar, Afilias, and ICANN.

The Conficker cabal assumed the botnet would be used for malevolent purposes. Due to the speed and breadth of the botnet's propagation, the responding actors wanted to stop further propagation of the botnet. Extreme time constraints arose when Conficker C first appeared in early March and threatened to undermine the previous response efforts of the community. The spectre of failure for the response efforts caused increased tension among them. The time pressure left little room to alter legal structures, e.g., to allow benevolent hacking of drones. This certainly also applied to possible changes in the governance or technical architecture of Internet addressing the attractiveness of the DNS system for botherders.

To respond to an existing and quickly propagating botnet, a number of possible strategies are at hand, with their applicability depending on a variety of factors. They range from eliminating the bots, preventing non-infected machines from becoming drones, severing the ties between bots and their command-and-control level ,to going after the perpetrators to reducing the value of bots for the perpetrators.[92]

An obvious, yet trigger-happy and mostly illegal approach to counter existing botnets is to eliminate existing bots. This can be achieved either by taking drones offline or removing the malware that turns a system into a botnet drone. Two very different organisational approaches can implement the technical measures necessary, either by existing administrative control within an organisation or remotely by a third party by targeted 'benevolent' hacking of known bots or by infiltrating the botnet and using it to propagate "white worms" (Hogben et al., 2011, p. 136).[93] A legal way to take out bots is by software updates provided by software vendors, but initiated by users. Updating operating systems or security software is daily business for AV companies and OS vendors. In the case of the Stormbot botnet in 2008, for example, the issuance of a software update by Microsoft helped to mitigate a botnet (Keizer 2008), but the botnet eventually regained strength.

[92] Similarly a report by the European Network and Information Security Agency, which outlines several "directions", "approaches", and "actions" against the "botnet threat" (Hogben, Plohmann, Gerhards-Padilla, & Leder, 2011, pp. 120-132). Another, albeit brief description of possible anti-botnet strategies can be found in Stanković & Simić, 2009.

[93] The case and proof of concept for mitigating botnets by hacking them has been made in Wicherski, Werner, Leder, & Schlösser, 2008.

A take-down procedure for all C&C domains and servers was not feasible with any of the previous approaches to fight an emerging botnet/malware threat. The bot-herders used HTTP rendezvous, first with 250 domains per day for each Conficker A and B and then 50,000 potential domains for Conficker C to pass new instructions to the bots. Never before was HTTP rendezvous taken to such a level. Even in the first months with just Conficker A and B, blocking communication between version A and B was a challenge, but some took up the gauntlet.

A second strategy against a rapidly propagating botnet is to harden non-infected systems targeted by the malware. To achieve this, security patches for exploited, vulnerable software components need to be developed and distributed. Likewise, updating malware signature files for traffic surveillance security products like AV software or intrusion detection systems (IDS) blocks uninfected machines from infectious data packages sent out by existing drones. Theoretically, this strategy would be sufficient if applied to every targeted system.

A third way to render a botnet ineffective is to take down the command-and-control structure. This approach had been very common in the heyday of IRC-based botnets with their central and therefore vulnerable architecture. These IRC-based botnets usually relied on command servers located at one or a few IP addresses. IP addresses are managed in a decentralised manner by five regional Internet registries (RIRs), each roughly covering one world region. These Internet authorities allocate the IP addresses to so-called Autonomous Systems, identifiable by their distinct Autonomous System Number (ASN) and managed by identifiable organisations such as corporations, universities or research institutions. The location of a specific server with an individual IP address can therefore usually be traced back with the help of the operator of an AS or its subordinated networks. Consequently, unless a server is located in a country without functioning law enforcement and hosted by a non-cooperative ISP, it can be taken down. A non-technical approach to taking out the command structure is to immobilize the bot-herder, for example with a traditional arrest by police forces.

Of all that we have listed, the fourth basic strategy — severing bots and the C&C structure — sounds the most unlikely. The Internet was built to ensure end-to-end communication between two nodes no matter what happened to some nodes previously located en route and in-between the two nodes. And yet, bots can be hindered from contacting their command-and-control servers as long as that line of communication includes resources that can with modest efforts be controlled by the defence side. The characteristic of attacking technique *HTTP rendezvous* is that bots try to reach their command servers under a specific URL. With *fast-flux*, a bot tries to connect to the command server using a specific URL; the DNS entry for that URL however is constantly altered, pointing to frequently alternating IP addresses. The bot-botherder communication for both *HTTP rendezvous* and *fast-flux* relies on DNS entries as envisaged by the botherder. Both techniques have turned control over entries in the domain name system into a critical resource for Internet security already in the years before Conficker. Taking down the actual command servers was no longer possible as the DNS system entry, quickly changed by the botherder, could theoretically point to any IP address. Supported by TLD operators, the defensive side could now sever the communication between bots and their C&C system by taking over or blocking the one or few domain names the bots tried to connect to. By severing bots and the C&C servers a botnet is rendered non-operational and useless, even when both bots and C&C systems are up and running.

Which of the aforementioned strategies fits best to the situation depends on a number of variables, among them time constraints, resources and the capabilities of the response side. When the persons who would become the core Conficker cabal had reached out to one another and discussed their options, they agreed on a response strategy called "defensive DNS". While the security and infrastructural-Internet industries were watching the emergence of Conficker, the initial response activities were performed by a small number of individuals, partly supported by corporate employers, partly self-supported.

Defensive DNS has been applied in response efforts before CWG. To counter the rise of the botnet Srizbi, security service provider FireEye had reverse engineered the malware's code and thereby figured out which URLs the botnet required to receive instructions from the commanding instance. Consequently, FireEye registered these domains to disrupt the communication channel between the botherders and the bots.[94] FireEye eventually gave up spending money on registering domain names, which allowed the botherders to regain control (Rendon Group 2011, p. 15). So, the response strategy was built on previous organisational and technological concepts. The Conficker malware was built on the shoulders of years of malware development. The organisation of the response likewise relied on existing paths of botnet and incident response.

[94] Rod Rasmussen, "Case studies in global criminal attacks: McColo & CheckFreea, presentation at ICANN 34, E-Crime and Abuse of the DNS Forum: Session 2, Criminal Attacks & Abuse Response Today, March 4, 2009, Mexico City,
http://mex.icann.org/files/meetings/mexico2009/ecrime-case-studies-04mar09-en.pdf

5.2.3 Informal security communities

The emergence of botnets in the early to mid-Noughties did not go unnoticed by security experts and the ICT industry. In 2003, Microsoft was living through its *annus horribilis*, "the worst in virus history" and a "nightmare" of an August. (Mikko Hyppönen, quoted in: Shnayerson 2004) That month, the Blaster worm took over a substantial populace of Internet-connected Windows machines, after a Polish hacker group spread the news about a serious vulnerability in Windows and Microsoft had immediately released a security fix. Understandably, Microsoft developed a dislike for malware exploiting wormable vulnerabilities. When in 2004 the authors of the Agobot botnet implemented an exploit that attacked the Windows' LSASS vulnerability that had previously been attacked by the notorious Sasser worm, "Agobot drew the wrath of mighty Microsoft" (Canavan 2005, p. 18). The author of the malware was eventually convicted after a cooperative effort by Microsoft and international law enforcement agencies (Canavan 2005, p. 18; Leyden 2014/2004).

In October 2004, Microsoft hosted the first International Botnet Task Force Conference, aimed at training law enforcement officers worldwide (Charney 2005, p. 6; Microsoft 2009a, pp. 23-28). In subsequent years, the biannual Task Force meetings gathered researchers, law enforcement officers and academics (Kaplan 2008). Microsoft also deepened relations to law enforcement and acquired law enforcement competency, e.g., by hiring former investigators from high tech crime teams at law enforcement agencies (Gotlieb 2011). The IBTF was not Microsoft's only forum to outreach and unite to other players in the field in the struggle against malware and cybercrime. Since the Slammer and Blaster worms in the early 2000s, Microsoft has initiated a torrent of three or four-letter-acronym security forums: "the Microsoft Security Response Alliance (MSRA)", "the Global Infrastructure Alliance for Internet Safety (GIAIS), the Microsoft Virus Initiative (MVI), the Virus Information Alliance (VIA), the Security Cooperation Program (SCP), … the Microsoft Security Support Alliance (MSSA)", and the "Microsoft Active Protection Program (MAPP)" (Microsoft 2009a, pp. 15-16).

Besides such corporate liaising, groups of like-minded security professionals convened at low-key conferences dedicated to the emerging botnet problem. In August 2006, the "Internet Security Operations and Intelligence" (ISOI) workshop, which was also dubbed "DA workshop" after the botnet-focussed Drones Army (DA) mailing list set up by Gadi Evron in 2004, brought together some of the individuals that were to play key roles during the forthcoming Conficker crisis.[95] The workshop gathered botnet-interested security professionals that were members of one of the security-related mailing lists OARC, NSP-SEC, FIRST, or were members of the Honeynet Project.[96] These early efforts helped to bring together experts with various backgrounds to address the upcoming cybercrime issue and coordinate responses against global incidents.

The analysis of malware in its various flavours plays an important role in almost any incident response case. Two examples of malware-oriented mailings lists are the Yet Another Security Mailing List (YASML) and Incidents & Insights (II). Not much is known about them in publicly available sources. YASML is apparently hosted by opensecnet.com[97] and grants access after a vetting process conducted by the list's steering group. The II list[98] is run by Ken Dunham, author of one of the first books on botnets (Dunham & Melnick, 2008). According to a presentation by John Kristoff — a former lecturer at DePaul University in Chicago and now researcher at the prolific security company Team Cymru — both mailing lists are used for multiple reasons. These include the sharing of samples of malware, along with methods of exploiting technical systems, requests for assistance for reverse-engineering malware, and requests to initiate collaboration with third-parties who are also members of the respective list (Kristoff 2010, p. 11). The Malicious Websites and Phishing (MWP) list was initiated in 2004 to pull together security researchers and professionals, AV vendors, security providers, and law enforcement. The list, set up by Gadi Evron just like the DA mailing list, was considered a break-through at that time and was in the vanguard of the by now widely prac-

[95] The workshops agenda is still available at http://isotf.org/isoi.html. Drones Army: Drone is a synonym for a bot in computer security circles.

[96] As of February 2014, the workshop's program is still online at http://isotf.org/isoi.html. Evron called the group "Internet Security Operations Task Force", its self-description reads: "ISOTF is an Internet community dedicated to improving the security of the Internet via operational sub-groups focusing on specific problems, strategically focused discussion groupe [sic; AS], and the biannual ISOI conferences." (http://isotf.org/?page_value=10) Further mailing list-based communities are described elsewhere in this manuscript, for example, in section 5.2.1, Boot-up of the response.

[97] According to an HBGary email leaked by Anonymous, a member of the CWG appears to be involved.

[98] The list's homepage is at http://npogroups.org/lists/info/ii.

ticed networked collaboration between various parties with a stake in the Internet security domain.[99]

Attacks are discovered, mitigating measures are implemented at the screens of network and system administrators, security analysts and other technical personnel. Malware specialists are indispensable to understand how a piece of malicious software acts. But it is not they who have the capability or authority to implement the technical measures required against ongoing attacks. This is where so-called operations security communities step in. Several communities have emerged, at times dedicated to distinct operational technical aspects of the Internet such DNS, or BGP (Interview 87).[100] After Conficker, Operational security communities with a focus on internet security facilitate cooperation among operational staff to mitigate or solve ongoing attacks or incidents. The term operations security — often synonymously used with *operational security* and abbreviated as OPSEC — describes the security of operations aimed at, unsurprisingly, increasing security. The term is frequently used in traditional security institutions. The U.S. Marine Corps has the following definition: "OPSEC is concerned with denying critical information about friendly forces to the enemy." (U.S. Marine Corps 2001, p. 33) It is a perpetual process that accompanies any operation that could be disturbed by foes. "Operations security (OPSEC) is concerned with identifying, controlling, and protecting the generally unclassified evidence that is associated with sensitive operations and activities." (U.S. Joint Chiefs of Staff 1997, p. v) This kind of thinking has found a place in the Internet security community, especially after Conficker. OpSecTrust, also dubbed as OPSEC-Trust or Operations Security Trust, was established after the Conficker incident to create a large pool of community-vetted security experts who could be drawn upon in subsequent incidents.[101]

Of crucial importance for ensuring the functionality of the Internet is the continuous interplay of those technical experts controlling the Internet's backbone; i.e., its main routes, crossroads, and traffic lights. It is their networks through which eventually any data packet on the Internet must travel to get from one end to the oth-

[99] The list's homepage is available at https://linuxbox.org/cgi-bin/mailman/listinfo/mwp. The list is access-restricted, it needs the list-owner's approval to be admitted to the list. Evron has discussed his mailing lists in an article he co-authored (Solomon & Evron, 2006).

[100] DNS-OARC is the "DNS Operations, Analysis, and Research Center", available at https://www.dns-oarc.net. For a wider picture of existing security communities cp. Greene 2012b, p. 13, 2012a, p. 9.

[101] The website of OpSecTrust is available at: https://ops-trust.net (last accessed in February 2014).

er.[102] These communities of backbone operators are usually virtual communities of mutually trusted individuals communicating via an access-restricted mailing list, combined with annual or bi-annual meetings in the physical world.

In the early 2000s, operators of Internet backbones found their data tubes flooded and jammed with traffic that originated from the then ubiquitous Windows worms. The quality of their services deteriorated, while the service level agreements with their clients remained in place. Backbone operators were incentivised to step in and stop such malicious traffic. They had to align their activities with their peer operators to achieve sustainable effects. In 2002 however, even the simplest resources needed to reach out to other network operators were lacking. Packet Clearing House filled the gap with its INOC-DBA, the Inter-Network Operations Center Dial-by-ASN. This VoiP-based service enables a network operator responsible for an Autonomous System Number to call a person from another ASN, which might currently be the source of malicious traffic for the first network. Currently, the hotline phone system connects about 1300 Internet network-operating organisations around the world (Interview 23).

The nsp-security (NSP-Sec) community goes a step beyond INOC-DBA's dial-up service. To effectively react to ongoing incidents, tools for easy one-to-many communications were necessary. The NSP-Sec community facilitates communication between operational staff from IP transit, large content and Internet service providers. The list allows them to share forensics and monitoring data, and to coordinate joint activities against ongoing security issues, for example by taking down systems or blocking traffic involved in a particular attack. The community's goal is to empower technical staff to jointly implement measures against attacks to ensure the overall health of the Internet and the quality of transit, content and service provider services. The community and its mailing list are focused on collaborative monitoring and mitigation and are not just an informal place to exchange information.[103] Its roughly 250 members are expected to act upon request by others on

[102] ICANN/IANA have assigned blocks of IP addresses to RIRs, RIRs have allocated blocks of these blocks to Autonomous Systems, each of which has its own unique number, an ASN. The networks of an ASN share the same IP prefix — e.g., 201.2. or 233. or 55.4.1 —, and all addresses on the Internet with this prefix belong to the same ASN. A single ASN can be operated by one or more network operators (Hawkins & Bates, 1996). As a fictitious example, the hypothetic ASN38xx83 has been assigned to RIPE NCC, the RIR responsible for the Western half of the Eurasian continent. ASN38xx83 is then technically operated by several ISPs.

[103] The sub-mailing-list NSP-SEC-Discussion allows for wider discussions. In addition to that and the main mailing list, NSP-Sec has three language-specific mailing lists for China, Japan and Brazil, and one for the Asian-Pacific region. The NSP-SEC-LEO group encompasses "NSP-SEC members, law enforcement officers, and legal teams working on cyber-criminal take downs" (Greene

Footnote continued on the next page.

the mailing list; i.e., an operations staff member would act upon the note of his peer operations staff working for a competing company (Interview, anon.).[104] NSP-Sec played an important role against the Slammer worm in 2002 (Greene 2012b, pp. 32-33) and contributed to the response to the Estonian cyberattacks (cp. section 4.1).

Another example of such a specialised operational community is DNS-OARC. The DNS Operations, Analysis, and Research Center brings together technical DNS experts and operators of large and important DNS servers including root and TLD name-servers, vendors and service providers of DNS technology, and DNS-related security providers.[105] DNS-OARC is a non-profit, funded by membership fees (Kristoff 2010). It provides much of the functionality that a DNS-CERT would have provided, if plans for such an organisations hadn't faltered (Interview 86).

All these communities, networks, and mailing lists were the platform that the Conficker Cabal and then the Conficker Working Group was built on. More than half a decade of conferencing, networking, and responding to Internet incidents, malware, and botnets have created the prerequisites for a networked response to the Conficker menace.

5.2.4 Structuring community-based response

The response to the MS08-068 vulnerability and then eventually the Conficker malware started at various loose ends, which were partly knit together in different existing mailing lists. Soon after the news of the malware broke, a dedicated mailing list for the response to this incident was created by members of the Shadowserver Foundation in mid-December (Bowden 2011, p. 97).[106] This mailing list constituted the communicational backbone of the response and helped to create an

2003). NSP-SEC could be seen as an exclusive club with the wider network operators community. The NANOG-mailing list of the North American Network Operators' Group is an, other that the security mailing lists discussed here, open discussion platform. Its archives are accessible under http://www.merit.edu/mail.archives/nanog.

[104] Cp. the mailing list's homepage, available at http://puck.nether.net/mailman/listinfo/nsp-security.

[105] Cp. homepage of DNS-OARC, available at https://www.dns-oarc.net/

[106] The Rendon report states that Shadowserver set up the Conficker mailing list only on January 28 (Rendon Group 2011, p. 17). This date is awkwardly late and contradicts other statements. In addition, the Rendon report appears to have quite a number of flaws when it comes to minor factual details.

awkward blending of a mailing-list-based community and a virtual ad-hoc organisation. In these early days, only persons that had known each other for years made it onto the mailing list.

Eventually, Conficker headlines spilled over from geeky mailing lists to IT-focused media and eventually to mainstream media. The Conficker response became interesting to more than just security experts because of the technological challenges. Marketing was a driving motivation luring in "PR flags" and others who "figured their company had interest or were assigned to get involved", one of the early members argues (Interview 38). The Cabal coordinated their activities and shared their views also on weekly conference calls, which regularly attracted some fifty persons (Interview 38). It would, however, be unfair to blame the increase of the Cabal membership entirely on selfish motivations within the IT security industry. The very technical nature of the attacks, the response strategy chosen and the technical and organisational givens of the DNS system required a large response community.

Such a rise of membership required organisational changes. The initial members of the Cabal pondered the options "throughout January" and eventually implemented "large-scale coordination" mechanisms by early February (Rendon Group 2011, p. 17). With eventually almost 500 persons involved in the response effort (Interview 48), the unstructured flow of communication and information became hardly manageable. The mailing list, with its torrent of uncategorized mails, became an impediment to effective collaboration (Interview 48).

> "The working group soon exploded to hundreds of people, it rapidly fell off the tracks and deteriorated pretty quickly. … There were guys who were doing the malware reversing, they would talk all the time. The TLD group doesn't need to know that, they don't care." (Interview 85)

Furthermore, communication with law enforcement, which became increasingly interested in the case even though it had little capacity or willingness to contribute, required a distinct mailing list. As a solution, seven subgroups and one core group were created: Core Management, Malware analysis, DNS registration, Sinkhole data, Remediation, Public Relations (Rendon Group 2011, p. 44). The activities related to the second to fourth groups are discussed in the subsequent sections. This restructuring allowed for more effective[107] collaboration on different subjects, but it also let the core group "regain control of the effort" (Interview 85). From

[107] This interpretation however is not unanimously shared within the community (Interview 47).

then on, contributors were vetted before they were placed in subgroup mailing lists (Interview 85).

A second way in which the organisation of the group matured was a change of wording. The somewhat ironic, self-mocking name Conficker Cabal was replaced by the reassuring, professional expression Conficker Working Group. With great fanfare, the CWG was announced publicly with the reassuring weight of all its 30 organisational members.[108] As of this writing, the CWG still exists.

> "So the core team, that dozen people, still does communicate, and we still have to make a good decision every year, do we keep collecting the data, that kind of thing." (Interview 38)

This leads us to the decisive question: what kind of security services were provided during the incident?

5.2.5 Malware and botnet analysis

Conficker analysis was not a routine job. The AV industry's core business is to sell AV software and subscriptions to the latest malware signatures. Back in 2008, the creation of a signature file was a labour-intensive handcraft (Egele, Scholte, Kirda, & Kruegel, 2008) that blended assembly line production with cutting-edge research. In 2012 the situation was not very different from 2008/9 — some 100,000 new pieces of malware, identified by unique MD5 hashes, ended up in the honeypots and shared signature databases of the AV industry, per week. Polymorphism of malware — think of a piece of malware that dresses up in hundreds of ways but basically does the same things — has led to skyrocketing numbers of seemingly new malware. The numbers of new malware files per day may be huge, but the risks they pose are probably significantly lower. Journalist and writer Ed Bott: "Counting every signature is an easy way to get to an impressively large number, but it isn't an accurate way to assess the current threat landscape." (2012) The comparison of two consecutive random months showed that Symantec's definition

[108] List of members include: 1and1, Afilias, AOL, Arbor Networks, Cisco, ESET, F-Secure, Facebook, Georgia Institute of Technology, Global Domains International, IBM-ISS, ICANN, Internet Storm Center, Internet Systems Consortium, IT-ISAC, Juniper, Kaspersky, McAfee, Microsoft, Neustar, NIC Chile, OpenDNS, SecureWorks, Shadowserver, Sophos, SRI International, Support Intelligence, Symantec, Team Cymru, Trend Micro, Verisign. Source: CWG homepage, available at confickerworkinggroup.org.

database showed only two "new name named entries" per day.[109] In a way, polymorphism and other techniques to disguise malware acts like a DDoS against the business models of the AV industry. Unsurprisingly, much research has been devoted to automate classification and malware analysis (Egele, Scholte, Kirda, & Kruegel, 2012).

The standard treatment for a new piece of malware has been behavioural analysis, a task performed by several hundreds, if not thousands of malware analysts worldwide on a daily basis (Interview 73). Researchers run the malware in a sandbox, a guarded system controlled by the researchers, for approximately "five minutes" (Interview 47) to see how it behaves and compare it with existing malware.[110] Getting hold of a malware sample can at times be difficult for independent security researchers. As a worm with a high proliferation, however, Conficker struck out into every network and researchers' honeypot (more on sharing of malware samples in the subsequent chapter).

Behavioural analysis of Conficker revealed its network traffic and connections. By simply playing with the date of Conficker's sandboxed host systems, researchers observed that the malware would contact various popular domains at certain dates. Exploring the future behaviour of the malware requires creating an environment for the malware in which it believes it actually runs at a future date. Conficker phones to popular websites to extract the timestamps from their responses. Therefore, analysts in the labs had to use proxy servers that imitated the behaviour of these websites. At a certain point, it becomes cheaper to reverse engineer a malware than to invest into an artificial sandboxed environment to make the malware believe it runs on a regularly infected machine.

Reverse engineering of the malware's binary code often is the only way for the responding side to understand the techniques of the botnets so that they can design

[109] The longer version of Bott's argument: Symantec's Virus definition lists for their AV software contains 17,5m entries; it roughly grows by 100k new signatures per week, identified by unique MD5 hashes (this was roughly backed by an interviewee who spoke of 150k new malware files per week); the list of new actual malware however is much shorter, only 213 new trojans, worms or viruses per week with some of them threatening old, unpatched Windows 2000 systems or delivering malware for already decapitated botnets.

[110] Slightly off-topic and with thick grains of salt, a rough estimate of the economics of malware analysis: According to the figures given above, behavioural analysis for 100k new malware would require 500k min computing time. With a week having 10,080 minutes, this require 500 machines for a super-quick automated 5-minutes analysis and probably for some times as many staff if this all still requires as much human intervention as in 2008. So, the bread-and-butter business of the AV industry is highly capital- and labour-intense.

appropriate mitigation techniques. Conficker was a complicated, challenging, and pressing malware. It consisted of several modules which to design and therefore to understand required a wide set of deep computing skills ranging from Windows' deep internals, to a wide set of malware techniques, peer-to-peer technologies to cutting-edge cryptography. One of the earliest aspects that the analysts tried to understand was the domain generation algorithm. By mid-December this and much of the functionality of Conficker had been reverse-engineered by various contributors from the AV industry, research institutes, and volunteering individuals. The most thorough analysis came from a team of researchers at the SRI International research institute (Porras et al., 2009/2009a, 2009b, 2009/2009b, 2009a), followed by the report from the Honeynet Project (Leder & Werner, 2009). The latter also contributed methods to inject and poison peer-to-peer tables to block communication via the P2P mechanisms implemented by version C. The encrypted communication channels between bots and the control layer of the botnet could not be cracked though. Some members of the Cabal thought about asking the NSA for help, but that idea was not widely welcomed by the community (Interview 47; Interview, anon.). Researchers also worked with Intrusion Detection System (IDS) vendors to create signatures for IDSs to find and block Conficker-initiated traffic (Interview 13). Members of the cabal were full of praise for the contributions of the researchers. Nevertheless, researchers were greeted with some scepticism as none of them had previously worked with core members and therefore had no trust relationship with the core group.

5.2.6 Defensive DNS

The principle of Defensive DNS is rather simple: The bots are prevented from reaching the botnet's command layer by ensuring that the requested domain names are resolved to IP addresses and systems that are controlled by the defence side. Organisationally, defensive DNS is more demanding, especially when there are not just a dozen or hundreds of domains, but 18.5 million domains per year as in the case of the Conficker family. The implementation of defensive DNS as it has been designed for Conficker and still is in place sounds almost easy. Thanks to the help of the reverse-engineered domain generation algorithm, the CWG creates a new list of domain names every year.[111] These lists are then sent to the TLD operators,

[111] Example of these lists is the "2012 Conficker Domain List", a 460 MB text file available on the homepage of the Conficker Working Group, http://www.confickerworkinggroup.org/domains/confickerdomains-abc.txt.zip. Entry per line is "2012-01-01 c arjcg.my", the date indicates when the bots would visit the domain to check for updates, the single letter indicates the Conficker version (a, b, c), and the last segment contains the

Footnote continued on the next page.

which then make sure that these addresses are routed to either dev0 or to a sinkhole server operated by the CWG. (China operates its own sinkhole servers, though.) This botnet telemetry has helped the Cabal/CWG to monitor and analyse the botnet and its expansion. No monetary transactions are required for this process to happen. ICANN and the domain name registries have waived all fees for Conficker-related domain registration or blocking. Things looked very different in the beginning.

In the early stages, a few individuals responded to the emerging threat on their own terms. When it became apparent that Conficker used HTTP rendezvous and defensive DNS was regarded as the response strategy of choice, the still small Cabal used a popular costless method to register domains called domain tasting. Until summer 2008, ICANN charged no fees from registrars when domain names were returned within a five-day grace period. Eventually, someone realized that owning a domain name for five days for free could be turned into a business. Given a particular characteristic of Conficker's HTTP rendezvous system — the malware-generated domain names were only used for one day — free domain tasting would have been a perfect fit for CWG's defensive DNS strategy. But when it transpired that more than 99.7 percent of all domain name registrations were speculative and returned within five days, ICANN and the DNS community decided to drastically reduce priceless domain-tasting in 2008 (ICANN 2009). This action came with an externality, for which some Cabal members, an AV company and a generous individual, had to literally pay for only a few months later. When the cabal had to acquire 2nd level domains en masse, these volunteers spent about 60 cents per domain name on average (Interview 21).

As a first means to redistribute this unfavourable distribution of the cost of security production, the cabal sought to pull in the operators of those TLDs that Conficker used in its domain generation algorithm. The cabal approached Verisign for .com and .net, Neustar for .biz, and Afilias for .org and .info (cp. Rendon Group 2011, p. 19). They all consented and representatives became part of the response endeavour. Neustar in addition demanded ICANN to waive the fees a registrar has to pay to ICANN per domain. And indeed, ICANN was eventually brought in to help reduce the financial strain for the cabal. With Conficker C, ICANN's participation was inevitable. Without them, the defensive DNS approach would hardly be sus-

domain that the bot would visit and which therefore is to be blocked. For any given day, it contains 50,500 entries, 250 for Conficker A and B each, and 50,000 for Conficker C, resulting in almost 18,5m per year.

tainable any longer, economically, but also organisationally. ICANN, too, had to enter new territories:

> "Conficker B was easy. Conficker C was a step into a minefield for ICANN." (Interview 86)

With no contractual or legislative leverage over country code TLDs, ICANN had to act persuasively and convince TLD operators that the measures suggested by the widely unknown CWG would better be implemented for the sake of the general Internet health. The CWG prepared clear and concise instructions for TLD operators, and ICANN representatives added their credentials to the Cabal's requests for support (Interview 86).

TLDs are operated by organisations with widely differing capabilities, cultures and environments. Experts estimate that the operation of a state-of-the-art TLD with adequate resilience against DDoS attacks and other security mechanism requires at least a $1m annual budget. But many smaller countries would operate their TLDs in a less ambitious way.

> "They are small organizations, sometimes in university, sometimes in businesses, sometimes in part of the government, where running the top level domain is just a secondary part of their day job. It's not what they do all day." (Interview 86)

Consequentially, not all TLDs are "highly advanced organizations" (Interview 86) — a fact that the CWG had to take into account. Not all of them had the time to develop their own automation mechanism to feed their DNS system with the blacklists generated by the reverse-engineered domain generation algorithm (Interview 21).

The domain registration is still ongoing, the CWG still blocks it year after year (Interview 21). It requires very little human intervention. The domain registration is highly automated. Only quality insurance, making sure that all the necessary domains actually are owned by the response community, requires human activity (Interview 48). Only domain collisions have and still do generate some work. At times, the Conficker's DGA creates a combination of letters that is used as a domain name. Then manual research is required to judge whether these domains are legitimate or need to be taken down (Rendon Group 2011, p. 18).

5.2.7 Sinkholing

To understand the activities of the bots and the size of the botnets, so called sinkholes were created. Sinkholes are central repositories storing traffic data sent from the bots to alleged HTTP rendezvous domains registered and controlled by the responding actors. These sinkhole databases contain information about Conficker-infected machines and their attempts to contact alleged rendezvous domains. But as these domains had been blocked and routed through sinkholes by TLD operators or Cabal members who bought them, the bots left behind traces in the sinkhole.

Operating sinkholes is less about mitigating and more about understanding a botnet problem. Sinkholed botnet telemetry delivers valuable data for mitigation purposes if one wanted to take down any bot, for example by remote disinfection. However, this path was not chosen by the response team as laid out in section 5.2.2. Especially in the beginning, data accumulated in sinkholes helped to understand the size and scope of the problem (2011, p. 17) and also the behaviour of a malware. It took a few weeks until Conficker was reverse engineered.

Initially, sinkholes were run independently by several person and parties in the core group. At its height, CWG members operated eight sinkholes. The Shadowserver Foundation, Richard Perlotto, Rick Wesson, and Chris Lee were operating private sinkholes; AOL, Microsoft, F-Secure, and Georgia Tech were among the organisations who had sinkholes running (Interview 47; Interview 48; Interview 21; Interview 13). The distributed, uncoordinated ad-hoc approach to defensive DNS and sinkholing led to some collisions.

> "So we would see some domains get registered, we were like, oh my god, is this the botmaster or is this just another researcher?" (Interview 21)

Thus, the community had to figure out who was running which sinkhole on which IP addresses. In addition and to enhance its situational knowledge, the cabal consolidated the data from different sinkholes in a partition of Amazon's cloud that was leased by one cabal member, co-administrated by others (Interview 21). Later on, that aggregation of sinkhole data was migrated to systems at Georgia Tech for the computing and bandwidth resources available there. In addition, a university was regarded as a fitting neutral ground for a valuable resource created by a community of persons and organisations with competing interests. The price tag of Amazon's services was certainly another motivation to move things over to Atlanta

(Interview 21). In its heyday, the cabal uploaded six to seven million unique IP addresses per day, Conficker's estimated populace (Interview 74).

Despite the centralisation of botnet telemetry, the sinkhole system still had elements of decentrality. To decrease the risks of DDoS against the central sinkhole, the systems were federated and distinct sinkholes were operated by distinct parties. Furthermore, telemetry traffic from bots to rendezvous domains was distributed among sinkholes using DNS techniques like GeoDNS and Round Robin (Interview 47; Interview 38; Interview 13). More sinkholes ensured that data ownership was distributed among several actors in the community. But it also made sinkhole operation more effective and resilient.

The CWGs sinkhole operation is still up and running on systems hosted by the Internet Systems Consortium (ISC), albeit on a slightly reduced level (Interview 48). In 2012, there are still some four sinkholes in operation (Interview 48). The reduced number of sinkholes over time has also lead to a reduced coverage of actual infections, but it would still capture around 75% of all infections (Interview 48). Some infected parties however blocked bots' traffic to sinkholes in order to avoid revealing information about their internal networks and infected machines. The number of IP addresses used by the CWG for their sinkholes were limited, usually in the same, say, Class B address range, and therefore predictable (Interview 38). Corporate IT departments, whose risk management procedures prescribe that their vulnerabilities can not be revealed to third parties, would consequentially block outgoing traffic to these address ranges by their firewalls. Therefore, known sinkholes within a fixed IP range might miss some infected machines and organisations in their statistics.

Akin to the case of the Estonian cyberattacks in section 5.1, section 5.2 added a narrative to the response to the Conficker botnet following the model of Internet security production in section 2.4.2. Apart from sanctioning, all security production processes have played major roles in the Conficker response. Monitoring, analysis, forensics, and mitigation have all been present, with many of them requiring the highest levels of technical expertise and solid operational execution. Malware and botnet analysis required state-of-the-art reverse-engineering and malware expertise. The Conficker sinkhole operation was probably among the first big-data Internet security projects. With Defensive DNS, finally, the Conficker response community implemented a highly specialised mitigation strategy, uniquely designed against the botnet.

5.3 Conclusion

This chapter has discussed the granular results of the response endeavours. Partly, this has been a research goal per se. The history of Internet security is only in its infancy, and the cases described here lacked a thorough account either entirely (the Estonian response) or in many aspects that are required to answer the remaining research questions (Conficker).

An essential aim of this chapter was to carve out distinct processes or products in the entire response effort. Due to the different technical design, quality and scope of the attacks, the response encompassed mostly different, but also some overlapping elements. A unifying element is the sharing of various kinds of information (more thereon later in section 6.2.3) in an ad-hoc, secure, trustworthy manner. Both response projects included malware analysis, but while marginal in the Estonian case, it was highly critical and extensive for the Conficker response. The Conficker case was dominated by malware and botnet analysis, and defensive DNS sinkholing. The Estonian response was less differentiated. Its main processes have been situational awareness, DDoS mitigation, information sharing, malware analysis, and policing.

One hypothesis is that those processes that do not require ownership over proprietary resources, that are entirely information-based and therefore have no hardware-related price tag, and that are more research-oriented are more prone to openness, distributiveness and socialness. That hypothesis will be scrutinized in meticulous detail in the following chapter.

One apparent finding here is that all and every incident response process depends on the contributions of the Internet security community. The security community — this vague ideational sum and umbrella construct of existing Internet security communities — and its values, norms and internal practices leave their marks on the way in which Internet security is produced and re-established after large-scale incidents. The achieved outcomes of the responses would not have been viable without these networks of mostly technical experts. Their information sharing, governed by a set of community norms, and collaboration allowed them to deliver services like malware analysis, defensive DNS, sinkholing, and DDoS mitigation. No other institutional arrangement had been available at the time of the cases to produce similar outcomes.

Both cases differ in many aspects. Other than in the Conficker case, the response communities have not defined their activities in this way, let alone adopted an organisational structure that followed such imaginary product lines. Arguably, the

Conficker Cabal/CWG had to reflect more thoroughly on their internal processes and activities than the more loosely coupled security and response communities in the Estonian case. The Conficker response stretched on over a longer period of time (five months vs. three-and-a-half weeks). Some in the Cabal had been involved in larger-scale botnet responses before. The members of the Cabal had more managerial experience, especially in terms of team-size.

Both cases touched on national security and even geopolitical issues, both with little actual impact, but both with a potential to shake up world politics substantially. Most striking however is how the response in both cases required ad-hoc collaboration with alien foreign experts. Arguably, the CWG depended on foreign support even more than the Estonians. While the Estonians could have somehow struggled through their ordeal, the CWGs response would have entirely broken apart without the support of every single TLD. In the Estonian case, a local national community reached out to European CERTs and then a subset of the global Internet security community. In the Conficker case, US-based persons involved in several global Internet security communities established a core group that later reached out to TLD operators worldwide. There were even some personal overlaps between both response endeavours: Finnish security and malware researchers; one US-based DDoS researcher; some two persons had at least indirect influence on the response activities.

It is not just this overlap of persons involved that highlights the importance of this informal, virtual Internet security community for incident-related Internet security production. It is the similarity of the norms, practices, and values in these communities world-wide that has defined the responses. These factors have arguably had a greater impact on the openness, distributiveness, and socialness of the response than the differences between the different security processes.

The next chapter follows the question of whether elements of peer production have been present in these efforts to re-establish a secure state after a security incident.

6 Social Dimension of Internet Security

Producing Internet security in the event of large-scale Internet security incidents means to re-establish the expected functionality of the Internet and the *status quo ante*. In the two cases analysed in the previous chapters, this somewhat intangible result is translated into a medley of processes and activities such as situational awareness, DDoS mitigation, malware and botnet analysis, defensive DNS sinkholing, policing, and information sharing.

Different theoretical frameworks allow us to view historical events, political struggles and social initiatives from different perspectives to discover new nuances and details that might have previously been ignored. Most analysis of security incidents and the responses to them focus on the effectiveness of the response and the potential damages that could be avoided by the response. The questions about the applicability and actual application of peer production, however, require a different perspective. Distributiveness, openness and socialness are the characteristics that matter here. Real-world applications of the peer production idea all incorporate these characteristics to a large extent.

The aim of this chapter is to analyse whether the responses to the two cases described in the previous two chapters have actually been distributed, open and social. The first section, Distributiveness, looks at the "peer" dimension of the response in both cases. Peer production requires widely distributed authority among relatively equal actors. This section therefore looks more closely at the topology of the response network, its governance model and the distribution and concepts of authority within the response endeavours. The second section, on Openness, depicts the accessibility of production platforms, input resources, and produced goods. It furthermore analyses the degrees of proprietarity of the results of the response efforts and the presence of the openness mentality among the response teams. The third section, Socialness, explores the social dimension of the response by analysing the motivations of the responding actors, the absence or presence of managerial and market influences, and finally the role of the state in the overall response. The chapter then concludes with a discussion of the results and an answer of the first research question on the existence of peer production in Internet security.

6.1 Distributiveness

Peer production refers to the collaborative creation of results by peers, a term that refers to a production group, community or network of persons with relatively homogenous authority status. As described in section 2.1.4, peer production is characterised by the "absence of a central hub or even decentral hubs with substantial control", and by a form of self-governance with dispersed authority within the network. Distributiveness is the term used to describe these characteristics. In peer production, power, authority and resources are distributed among the actors in the production network; the topology of the network is distributed, not centralised and not controlled by a single entity.[112]

If attacks on the Estonian infrastructure in 2007 had been exclusively physical and conducted by local demonstrators in several cities, the response would likely have come from a precisely organised network of national, regional and local state-operated security forces, police units, intelligence, and possibly military units, all presumably orchestrated by a national crisis committee with defined chains of command. If the attacks had been conducted by, say, Russian military forces, the response would have likely come from the alliance of NATO members, likely centrally orchestrated by a NATO crisis committee with the US pulling the decisive strings. But what are relations like between the actors that participated in the discussed responses? This subsection explores the distribution of authority among the

[112] One could argue here whether the best attribute to describe peer production is "decentralized", "radically decentralised" or "distributed". Benkler uses all concepts, Bauwens primarily the latter (cp. section 2.1.3). Referring to Galloway's influential book Protocol (2004), Bauwens (2005) argues that distributiveness would be the appropriate term. In a distributed network, the network topology may contain several hubs, but those could be routed around. In a decentral configuration, however, such hubs are not only an essential element, but they also cannot be routed around. In technical networks, such a differentiation is probably easier and a binary question. In networks comprising social actors, this is a more challenging task, especially when one takes the possibilities of social changes into account. Furthermore, while a technical decentral network may break when one decisive hub is taken out, social networks may still find ways to route around a gap in an otherwise decentral network.

response actors, the modes of governance and the topology of the response networks.

There is an underlying reason why questions of distributiveness in network topology, and centrality or federation in polity design are so relevant in political debates. These variables are linked to and are partly even synonymous with distribution of power and authority. The word *peer* refers to the, as Bauwens calls it, *equipotentiality* of the actors involved in peer production. In peer production, participants roughly have the same capabilities to influence the course of action of the response efforts, shape the behaviour of others involved, enforce social or regulatory norms, or sanction defection of others. Authority, let alone its distribution in society, is rarely visible. Sovereignty, a conceptual sibling of authority and considered as "absolute and perpetual power" (Bodin 1576/1955, pp. 25,28), only reveals itself in exceptional times; then the sovereign becomes visible and emerges as "he who decides on the exception" (Schmitt 1922/1985, p. 5).[113]

The relationship between production communities and authority has three dimensions, describing internal, outbound, and inbound authority. Internal authority describes the distribution of authority within a production community, the relationships between the members of a community. Outbound authority refers to the authority the production community exerts on non-members. Inbound authority describes forces that are not part of the community, but nevertheless shape the behaviour of participants of the production process. For the analysis of the presence of peer production in incident response networks, the internal dimension is decisive. Before discussing the empirics of distributiveness in the two cases, the subsequent section sums up some findings on internal authority in existing literature on open source communities.

6.1.1 Internal authority in peer production

Early evangelists of Internet culture and open source communities claimed that these communities were mostly egalitarian. With kings being defied, everybody had the right to contribute on equal terms.[114] Nevertheless, researchers of open source software or other types of peer production communities soon discovered elements of hierarchies within these communities and therefore the existence of

[113] Bodin and Schmitt are cited in Suganami 2007, pp. 513,516.

[114] "We reject kings, presidents and voting. We believe in rough consensus and running code." (Clark 1992)

authority therein (Weber 2004). "Whenever a group is focused on shared instrumental activity (rather than mere coexistence or expressive activity), effective authority is essential to success — to define direction, to allocate resources, and to resolve disputes." (Adler & Heckscher, 2006, p. 59) Such groups consequentially need to develop some kind of internal authority. What differentiates social groupings is how they organise the exertion of authority and which elements this authority addresses. Peer governance, which refers to "the way that peer production is organised" (Kostakis 2010), is a "bottom-up mode of participative decision-making where decisions are taken through the unconstrained engagement and free cooperation of producers" (Kostakis 2011, p. 151). For Michel Bauwens, who apparently coined the term, peer governance ensures the "creation of value" by peer production and preserves the preconditions of peer production; it "manage[s] this process and peer property to protect common value from private appropriation" (2012d).

In an empirical study of authorities in OSS projects, George Dafermos analyses the FreeBSD project and the changes of its governance model and its forms of authority over the years (Dafermos 2012, pp. 3-5). To cope with the increasing number of contributors, the governance path chosen by the FreeBSD community was not to manage the committers, and allocate them to prescribed tasks, just like an ordinary company would do. After all, self-allocation to tasks is one of the characteristics of peer production. Instead, Dafermos states the FreeBSD community democratised itself and replaced its previously self-installed, meritocratic governance board by an elected, time-limited version of the latter. Most importantly, the community harmonised the way contributors approach their tasks and communicate among themselves by a mentoring system that would guide talented coders to the inner circle of approved committers (cp. Figure 6.1). In addition, the introduction of nightly builds from FreeBSD's development branch ensured that flawed code is immediately detected (2012, p. 3). Dafermos argues that the emergence of this governance model can only be explained with the normative stance of the persons involved in the FreeBSD project. Given the hacker ethic and its focus on individual freedom and hatred of hierarchies, a central layer to control and organise contributors would likely have been a hit to contributors' motivations and lead to a decline in the overall effectiveness of the community, instead of increasing it (2012, pp. 5-9).

The implication is that for an adjustment of a governance system to be successful, one must be aware of the existing culture, norms, and values of the community. A traditional, hierarchical governance approach would likely have stalled the community; harmonizing input and nightly builds have made the community thrive. "The form of authority is closely related to the values of the community in which it operates. It can sustain itself only when it fits within these communal definitions of legitimacy." (Adler & Heckscher, 2006, p. 61)

6.1.2 Rough topology of the response networks

The analysis of networks in general and in the domain of security policies in particular is a complex undertaking. Entire doctoral theses have been devoted to the challenge of exploring policy networks and their impact on Internet security governance (e.g.: Kuerbis 2011). A similarly in-depth analysis of response networks would arguably require a different basis for this research, involving quantitative data and the application of social network methodologies (Ward, Stovel, & Sacks, 2011; Hafner-Burton, Kahler, & Montgomery, 2009). Appropriate data for such a quantitative approach would be archives of mailing lists of Internet security communities, for example, which so far have not been accessible for non-members. The aims for this subsection necessarily are more modest, more qualitative and interpretative in their approach. It will sketch out the nature of the response networks and highlight facts that either support or contradict the idea of a distributed response with distributed authority.

The Estonian response comprised various actors in Estonia and abroad. Noteworthy actors from Estonia have been CERT EE, national ISPs and telecom providers, some members from national intelligence and counterintelligence. The technical response was accompanied by political reactions and support from various government branches. The Estonian response network was accompanied by European and international CERTs, operational Internet security and malware analysis communities. These standing communities operated during the incident in their usual fashion. The international dimension of the response hence bears many of the characteristics of the communities that are at the centre of this thesis.

The Estonian response network has features of a dense network with interconnections between many players. There is low centrality in terms of who knows and interconnects with whom. An inner circle of security experts evolved as the fat nodes in the networks, primarily the security staff of domestic ISPs and banks, as they operated the largest and most relevant networks. Within this network, CERT EE became the decisive hub to distribute relevant information to relevant parties and coordinate some activities. It could only just fulfil this function by distributing tasks and functions to other CERTs and by neglecting functions like supporting foreign law enforcement agencies or documenting the incident. In a way, the response network partly routed around the congested CERT EE, which nevertheless remained a central unit in the response efforts.

In the Conficker case, categorising the network responding to the incident is somewhat tricky. The most important part of the response organisation has certainly been held by the Conficker cabal and then later by the core group of the CWG, a group of less than a dozen people. That group taken as an entity was a super-node of central importance for the entire response endeavour. The group was inevitable, unavoidable, required, and impossible to be routed around once it had been established. On the member-level, a number of actors certainly had central roles, i.e., it would have been hard to ignore or replace them. The response was comprised not only of the CWG, but of countless other actors as well. As an example: The entire response strategy built on the voluntary cooperation of top-level registry operators. Any of the latter could have opted to play the role of a weakest link that breaks the security chain of the response strategy.

The importance of the response network topology for Internet security governance has been highlighted by post-incident debates. The non-hierarchal nature of the Estonian response networks has been stressed by post-incident debates. The response successfully mitigated the attacks in Estonia; nevertheless, debates about the appropriate design of Internet security provisioning institutions arose. Policy circles seemingly favoured more traditional designs, and eventually succeeded in turning CERT EE's hosting organisation RIA into an *amet*, a national office with authoritative capabilities in future response campaigns. Members of the technical community opined that the importance of national CERTs should be reduced to increase the resilience of the response network. The successful attacks on the Finnish CERT had shown that the functioning of CERTs during a response is not a given. In fact, any serious attacker likely evaluates the organisational and technical design of the response side and goes after weak and decisive nodes. The resilience of a distributed network of peers contrasts with the authoritativeness of a Weberian bureaucracy (Related: Kostakis 2011; Kreiss, Finn, & Turner, 2011).

The Conficker case is an instructing example of the importance of the centrality or distributiveness of the collection, aggregation and analysis of security data. Such data is used to attain informational awareness, understand the problem or monitor the effectiveness of countermeasures. To mitigate Conficker botnet telemetry data was inevitable. The distributiveness or centrality of data, information, and knowledge, which is necessary to understand and eventually mitigate a problem, is paramount for the distributiveness of authority. The backend databases of the cabal's sinkholes comprise billions of datasets telling which Conficker bot attempted to connect to which top-level domain. Malware telemetry naturally is created in a distributed fashion by countless bots and needs to be collected either in a similarly distributed way or at least at many decentral gateways or passages of the Internet to get a coherent picture of a botnet's activity. While the sources of malware telemetry collection are necessarily distributed or at least decentral, the locus of data aggregation and analysis is likely central or at least decentral but with only a few nodes.

Whether control of such data and analysis resides with one, some, or a larger number of autonomous actors has been a revisited topic in the recent history of Internet security governance debates. In a way, it merely perpetuates the endless but unavoidable debate between federated and centralized political authority in any society. The landmark Markle reports, created in the aftermath of the 9/11 attacks, called for the utilisation of the full range of existing ICT capacities to pre-empt potential future terrorist plots (Markle Foundation - Task Force on National Security in the Information Age 2002, 2003). The authors foresaw potential problems for democratic control, legitimacy and privacy, and hence argued in favour of a federated monitoring system. Federation was seen as a barrier against a central information hub which would gather too much power, as it would give the ability to hold and combine tremendous amounts of data into the hands of a few. Many of the authors have held senior positions in the security complex since the reports were published (Schmidt 2013c). Nevertheless, the systems that have actually been implemented are anything but federated as the Snowden leaks revealed, but rather smell like an attempt to build a system of systems.

This debate not only pertains to the levels of national security politics or even geopolitical interests, but also to the relationship between users of information technology and ICT security providers. To give an example: In a malware monitoring system as distributed as possible, users have the ability to use any local monitoring components of their liking that communicate to one or several decentral data repositories. A system designed at the Internet Systems Corporation bears some of these characteristics. According to its mastermind and CWG member Paul Vixie, sensors installed at the discretion of users span a “cooperative, open sensor network” based on open source software and protocols. Sensor operators can then decide whether they want to use such sensor data by themselves or share it with other parties (Interview 38). This model contrasts with the model of central, proprietary networks favoured by other commercial outlets, be it AV companies, OS vendors or other security providers. Proprietary monitoring and security awareness systems usually consist of networks of distributed, yet proprietary sensors communicating with central systems that hold much of the intelligence. Sensors receive the latest information on current malware and detection techniques from their respective central system. In return, sensors send information about the situation in their domain to their central unit that is controlled by one organisation. This results in an ecosphere of several partly competing, partly overlapping, partly complementary monitoring and security systems. The lack of integration of these technical systems is partly made up by information exchange in Internet security communities. This is one of the community's core features: It provides a common platform to share distributed bits of data, leading to enhanced situational awareness and understanding; and thereby enabling distributed, yet coordinated responses despite the lack of central authority.

To summarise, the topology of the response has been rather similarly hybrid in both cases. A small group coordinated the overall response; it cooperated with a small or higher number of inevitable contributors (Estonian ISPs, TLD operators); and it was supported by an undefined number of distributed supporters.

6.1.3 Exerting internal authority

Internal authority is exerted in a variety of ways in groups and communities that respond to an ongoing Internet security incident. At times, authority is exerted in an obvious manner; at times, it requires the interpretation of singular occurrences during the response. Some displays of authority are similar to what can be observed in open source communities, other elements of authority seem unique and tailored to the characteristics of the security communities and security regulation. List ownership, access control, and internal hierarchies by nested communities provide opportunities to exert authority within the community. Examples of enforcement of norms, or the lack thereof, allow for some deductions about the distribution of authority within the community. Despite these indications for internal authority, the degree still appears to be modest, but somewhat more pronounced than in open source communities.

The response activities have been organised mainly via mailing lists and list-based communities. Those who have administrative control over these mailing lists regulate access to these communities, at least in theory. Some of the commonly used mailing lists in the community are controlled by one or few individual persons.

> "[H]e who operates such a list has pretty much influence as he can decide over who is getting on the list, what is the list's policy, i.e., how should people behave on the list, what happens with information shared on the list, which rules apply…." (Interview 37)

The list owners' ability to exert influence is shared by another person involved in the Conficker response:

> "They set up the mailing list, and there was a turn of control. Whoever sets up the mailing list first gets a lot of control. They own a big part of

> it. … they control information, and they control who has access to that information." (Interview 47)

Related to list ownership is the question of access to the community. The most obvious form of internal authority in security communities is their strict access policy, which is based on trust, resources, and skills. Depending on the policies of the respective community, prospects are vetted by an individual community organizer, a community board or in a collaborative effort by the entire community (more on that in section 6.2 later in this chapter). The type of control of community access indicates which type of authority is exerted within a community, be it charismatic leadership, meritocratic leadership by a self-imposed community principal or a group, or more democratic forms of governance, e.g., by collaborative vetting. In the Estonia case, initial access to the response teams relied on previous projects related to Estonian online elections and informal collaboration on Internet security issues. In the Conficker case, access to the core team relied on previous collaboration and a contributor's access to urgently required resources. In both cases however, response groups rose quickly in numbers, many persons were admitted and vetting was stripped down to minimal levels.

The Conficker Working Group has introduced elements of traditional organisational hierarchies and management layers into the response endeavour. New informational walls were erected that hindered the free flow of information among community members in different subgroups. Comparable to that is the idea of nested communities: a response group could invite other, less known and trusted, but necessary persons to participate in the response endeavour, without granting full access to community resources. Such hierarchisation is a manifestation of Benkler's assumption that peer production communities might have to refer to traditional hierarchies to integrate the contributors' deliveries to a cohesive product. In the Estonian response, internal hierarchal differentiations were not established. The overall internal governance was less sophisticated and more basic.

In a security production network, different approaches to internal authority are feasible. The decisive factors are the form of the deciding body and the bindingness of its decisions. As already mentioned earlier in this chapter, the deciding body could be an individual, a subgroup of the community or the entire community. Regarding bindingness, community members can be free to ignore decisions with no, little, some, or substantial repercussions for themselves. The default degree of formal authority of the network over its members is rather low. In the Estonian case, one characteristic of the collaborative approach was the somewhat central role of CERT EE within the network of actors. Nevertheless, even CERT EE lacked authoritativeness over any other actor. It could not "enforce its recommendations on all parties involved" (Evron 2008b). This does not come as a surprise given the wide distribution of ownership of those components, of which the 'Estonian' Internet consist, and how IT infrastructures are operated. Nevertheless, the Estonian response was "largely successful" (Duffy Marsan 2007) and executed "in a very proficient manner" (Ashmore 2009a, p. 6).

In the Conficker case, a similar absence of central authoritativeness can be observed. The response network in a wider sense also included nearly all TLD operators. (In the widest sense, the response networks included every ICT department that cleaned their systems and installed security updates or otherwise contributed to the response.) In the narrow sense, the response network is equal to the CWG. The core team of the cabal certainly had a strong influence over the entire design and execution of the response. But it had little leverage to enforce their decisions in the wider response network. Consequentially, it took persuasion and convincing arguments to get all TLDs on board. Only after Conficker did ICANN add provisions to contracts with registries that require the latter to perform certain emergency DNS security measures if requested.

A remarkable feature of both response efforts was that they were mostly driven by communities of technical experts with relatively equally distributed authority among the response participants, and little interference from commanding hierarchies or monetarily luring markets. The usual suspects to seize the position of a dominant authority, large international corporations or established political organisations, played a reduced role at best. If a crisis reveals "he who decides on the exception" (Carl Schmitt), it was the communities of technical experts and their reasoning derived from technical necessities and pragmatism that influenced the behaviour of traditional political and economic actors.

To contain the Conficker botnet, many influential actors had to ignore established procedures.

> "We got ICANN to break the rules, Verisign to break the rules ... everybody, and they all hated it. The countries that came in…, they broke all their own rules." (Interview 47)

ICANN and registries had to waive registration fees to make the response strategy economically feasible, registries world-wide had to trust the lists coming from the CWG and use them to block domains, some of them had already been registered by ordinary customers. Organisations shared information to an extent they had not done before. Data that revealed the vulnerability of computer systems was exchanged between politically conflicting parties. Largely, the CWG did not authoritatively impose their decisions onto others, but had to convince others to get onto their boat.

The situation in Estonia was similar and also required the cooperation of many actors in Estonia and worldwide. The symbolic sovereign of the moment was Estonia's technical community. During the weeks of the attacks, Estonian politicians, international journalists and diplomats were received in audience by the Estonian CERT. Even most senior national politicians travelled down from Tallinn's picturesque government district located on a plateau in the historic centre, to the more mundane district in which the RIA and CERT EE are located. After the incidents had been mitigated, the organisation and institutions of the response side have been altered by a number of initiatives driven forward by policy makers and from within corporate headquarters. In Estonia, the technical community has been institutionally intertwined with public authorities. Following the Conficker response, police high tech teams and large companies have gained a more prominent role in subsequent botnet responses. The changes may well have been driven by efficiency or security considerations. The institutional development after the incidents can however also be interpreted as an attempt to redistribute authority from loose, barely controlled technical communities to established organisations.

> "[It; the Conficker response] was a unique moment, because there weren't any rules, there was a great deal of stress and people wanted to solve a very perceived threat. And so people made decisions on the spot, and I don't think that that's gonna happen again." (Interview 47)

The rule of law and corporate regulations arguably replace the peer rule of technical reason and pragmatism that characterised the responses to the two incidents.

6.1.4 Deviation and internal sanctioning

Authority refers to the ability to determine the legal scope of action for others and thereby shape the behaviour of other actors. During the Conficker response, there have been cases in which a majority of actors built up pressure to assure certain behaviour or to sanction unwanted behaviour. Remarkably, no such occurrence has been mentioned for the Estonian case.

In his book *Worm*, Mark Bowden dedicated a large part of the eighth chapter to a conflict that arose within the core of the Conficker response team.[115] One of the initiators of the response activities unilaterally decided to pass infection data collected by his and the Cabal's sinkholes to a contact in the China Network Information Center (CNNIC), the Chinese public authority responsible for and operatoperating the .cn registry. CNNIC could then use this data to identify infected machines to clean them and reduce the huge share of Chinese bots in the Conficker botnet. The sharer's approach led to some distrust among other persons of the core group. His critics argued that such a decision about how to deal with data collected by the group should be made within and by the group. Some of his cabal peers accused him of using the Conficker malaise as an opportunity to get his emerging Internet security intelligence business rolling and sell data generated by the cabal's sinkholes to the Chinese. It would have been a plausible move after all. The data-sharer had spent thousands of his dollars to register Conficker rendezvous .cn domain names (2011, p. 147) (and tens of thousands for domain names at other TLDs). Collaboration of CNNIC and ICANN was not home and dry yet, and, after all, he had collected much of the infection data personally with his private sinkhole systems running on Amazon's cloud systems, which he had rented (Interview 47). He refutes allegations that he sold data, but does not dispute that he shared data (Interview 47).[116] The argument is supported by the fact that he was not ousted though and continued to contribute to the response efforts (Bowden 2011, p. 155).

> "[H]e's a bit of a maverick and just goes off and does stuff…. Some people have less tolerance for that than I do." (Interview 86)

When the CWG introduced a new organisational structure with a core team and a number of subgroups, he would no longer find himself on the "The Core Mailing List", the mailing list of the CWGs inner group (Interview 47). The group demoted one of its founding members. Despite that the data-sharer had spent almost US$100,000 on domain names to secure the cabal's response strategy before

[115] Bowden 2011, pp. 143-155. The remainder of this paragraph is based on Bowden's account, backed up by my own findings.

[116] It is not clear though which data was shared, a general sinkhole dump or those entries related to ASNs under CNNIC's auspices. Neither is clear when the data was shared. The data-sharer mentions that he flew out to an ICANN meeting in Mexico to get in contact with a Chinese person. ICANN held its 34th international public meeting in Mexico City between March 1 and 6. However, the meeting in Atlanta when ICANN agreed to reach out to China was in early February (Bowden 2011, p. 137). Apparently, a first contact was established already in the night after the meeting (Interview 86).

ICANN and the registries world-wide were on board, the cabal hasn't given him more leeway and make do with giving him a yellow card.

Whether it is a clash of personalities, a clash of a member with agreed community norms, or a disagreement about the existence of non-written, implicit community norms at the heart of the issue here, is of minor relevance when it comes to the question of how and to what degree internal authority was exerted in the response community. With access to the community as the scarce resource, informal authority can easily be exerted by revoking or reducing access to community assets. These means are not available in traditional open source communities as they use entirely open mailing lists from which by definition no one can be excluded. In an access-restricted configuration predominant in security communities, the threat of exclusion likely is a convincing argument for community members to live after the norms of the community. It is at least theoretically conceivable to exclude a disliked person by carefully framing their personality or actions as untrustworthy.

A second instance of internal authoritativeness in the Conficker response community occurred when a small group of independent security researchers were about to publish details about the botnet. The group of security researchers were arguably not thoroughly familiar with the disclosure practices of the Internet security community at that time. Openly publishing technical details on botnets in the wild would have, members of the CWG argue, created the usual problems for the response group that come with an open disclosure. Full disclosure would have informed the botherders about the current state of knowledge and capabilities of the response side. The botnet authors could then, according to the mainstream view within the CWG, adapt their malware accordingly to circumvent defence mechanisms by the response side. So far, uninvolved third parties could jump on the bandwagon and exploit the botnet for their own malicious purposes or strengthen the resilience of future botnets. Influential actors of the response side reacted wildly to this perceived threat, and were pondering whether to engage international police forces to confiscate the researchers' computers (Interview, anon.). The motives for this sharp response are not clear, but it has been argued that the Conficker cabal at that time was insecure about the feasibility of their response, the motives and future steps of the botherders, and signs of panic had emerged within the group.

Again, whether it was panic, momentary overreach, or sober judgement that led to the application of decisive authority is of secondary relevance. With regards to the question of the distribution of internal authority on the response side, the decisive point is that individual behaviour could plausibly be influenced in the way described. Furthermore, efforts that do not follow certain rules — here, disclosure rules — can be sanctioned. Security researchers at times happen to operate in legal grey areas, when they share sensitive data that might at times collude with privacy prescriptions or experiment with new defensive techniques that collide with non-hacking acts or other regulation that lags behind the latest developments in information security and insecurity. These facts can be used by actors with close links to LEAs to persuade other community members to act in certain ways. These factors create a structural advantage for community members with a strong legal department behind them over inexperienced individual players. In the case described above, this may have helped to enforce proven community norms.

The findings in this and the previous section suggest that whenever authoritativeness is exerted, the means of a community member's organisation is at work.

6.1.5 Decentrality and distributiveness

The response to both incidents definitely followed a networked approach. Whether it was more of a decentralised or a distributed effort is somewhat unclear.

These findings suggest that differentiating network topology by describing the role of hubs in these networks — hubs as optional and routable in distributed networks vs. required and impossible-to-be-routed-around in decentral networks — does not correspond to real-world social phenomena like the Estonian response. A decentral hub that is required and impossible-to-be-routed-around can be partly replaced by other nodes in the network. CERT EE's core functions in an incident response — gathering, analysing, and disseminating information — were inevitable for the response. But they were taken over ad-hoc by other nodes when CERT EE reached the limits. Consequentially, the above-mentioned differentiations between decentrality and distributiveness are insufficient to describe social realities. A node is only indispensible if other nodes cannot take over its decisive functions ad hoc. For real-life social systems this question is difficult to answer ex-ante when knowledge about a network is incomplete, a factor which might have boosted the popularity of cyber-manoeuvres in recent years. From a political perspective, the decisive criteria is whether a single or few actors can amass critical control over resources required to respond to an incident in an institutional environment that allows them to leverage this control over resources into influence over the design, practices and strategies of the entire incident response.

In neither of the two empirical cases have a few single nodes dominated the response. No actor emerged as the uncontested principal among others. Amorphous networks of actors called the shots. In both cases, the response was organised along technical requirements, status quo and adoption of security technology, existing institutions like local and global Internet security communities. Unsurprisingly it is apparent that certain functions or processes of security production have been in the hands of established institutions. Sanctioning and policing activities were in the hands of law enforcement and intelligence organisations, which have mostly monopolized these societal functions per national territory. This exclusivity combined with their surprising inability or reluctance to act led to a situation where these processes were not thoroughly covered, especially in the Conficker response. However, given that Conficker and, to a lesser extent, Estonia 2007 had global components, territorial monopoly is the equivalent of functional decentrality.

Functions like malware analysis were widely distributed and performed by AV companies, research institutions and independent individuals alike. Unlike most functions, malware analysis does not require ownership or control of network infrastructure, or a specific position in a private organisation or a public authority. Malware analysis first and foremost requires capturing a sample of the malware, high technical expertise, and years of practice. With the existence of open-source honeypot software, capturing a sample of a widely distributed malware takes only moderate effort, and the analysis itself only requires affordable machines and software.

Other processes are distributed as the resources for the response are distributed. In order to get a country-wide or even global picture about ongoing DDoS attacks, emerging botnets and malware proliferation, local information had to be gathered, consolidated and analysed. The lack of decentral monitoring capabilities at important Estonian Internet gateways made it necessary to gather data on ongoing attacks from distributed sources. To respond to an ongoing DDoS attack or botnet proliferation, organisational ICT systems need to be reconfigured or adjusted. Consequentially, security production processes like monitoring and mitigation required access to secured systems, controlled by respective organisations and their administrators. Given the structure of the Internet, all the components that require reconfiguration are privately owned, so access to those systems is limited to a few operators. In recent years, the exclusivity of administrative control over infected, private machines has been slightly reduced, e.g., by police forces hacking into and sanitising malware-infected machines, or by establishing public authorities with authoritative command in times of crisis (cp. section 6.3).

The distributiveness of administrative authority of ICT systems combined with the oft-lamented lack of incentives for organisations to harden their vulnerable systems forced the CWG to go for the DNS-based response strategy. One decisive component of the Conficker response approach, defensive DNS, required the cooperation of a plurality of TLD operators. While their task was pretty small and simple

— receiving a list of domain names and blocking or registering them — they were required and impossible-to-be-routed-around nodes in the response network. The same holds true for the largest Estonian ISPs, which were connected to the response network via their respective technical operations team members. With most of the inbound traffic passing through their networks, they were in a unique position to drop bad packets before they could congest the tubes to their clients' networks or overburden their systems. Therefore, the contribution of at least the majority of the largest Estonian ISPs was crucial for the response. Working with every attacked organisation was neither feasible, nor necessary. The gatekeeper role of ISPs allowed for a decentral mitigation process within Estonia.

Equipotentiality is a defining aspect of the relationship between peers. It involves internal governance mechanisms that don't discriminate between members. In the Estonian case, the organisational design was pretty simple and straightforward: a few meetings, an Estonian mailing list, later on international mailing lists, an Estonian IRC chat room. Post-incident complaints about the lack of an authoritative CERT EE response support the thesis of equipotentiality in the Estonian response. In the Conficker response, the initial cabal was a prototype of peer governance. With the formation of the CWG, however, the founding cabal members promoted themselves to a new managing level. The core group then oversaw newly added functional sub-communities, which by default operated separately from one another. An introduction of hierarchical levels has been done in traditional peer production and open source projects before; Linux kernel and especially the somewhat bureaucratized Wikipedia are examples of hierarchisation and internal authority in open source projects.

The question of authority overlaps with other analytical concepts that are described later in this chapter. Questions of inclusion and exclusion are discussed in the following section; the role of the state is touched in the section 6.3.4; norms of the community and relations between community members and employing corporations are described in section 6.3.

6.2 Openness

Openness has been a central concept and a catchall term for all goodness of the past Internet decade. It describes the compatibility of technical standards, the general reachability of websites, and the accessibility of democratic institutions. It is a descriptive term, just as it has become a normative goal, if not a panacea to so many societal ills. Openness is also the founding term of peer production's conceptual twin, open source production. This section therefore addresses the core of this study, which is driven by the puzzle of whether security production is reconcilable with openness.

Openness refers to several characteristics when applied to the context of Internet security production. As described in section 2.3, openness refers to the accessibility of the production platform, input resources and intermediary goods, internal transparency about activities and contributions, and the mental disposition towards openness. Openness also describes the accessibility and modifiability results, a commons-characteristic of the produced goods. Furthermore, the term refers to whether produced goods — as opposed to input resources and intermediary goods as mentioned before — are accessible, reusable and modifiable outside market/hierarchy-exchange frameworks, and whether they are accessible to non-producers. Last but not least, a truly open production project is also forkable.

This subchapter discusses these aspects of the concept of openness and analyses whether and to what extent Internet security production incorporated the principle of openness in the cases portrayed in this study.

6.2.1 Accessibility of the production platform

The production platform is where the workbenches are, the coffee bars where one meets for a chat and talks about current issues at work; they are the hacking clubs where one hangs out and codes. In the physical world and as long as printers that can create any type of connections between atoms are yet to be invented, producing requires people coming together in one space. In the industrial world, many such production facilities are locked down, hidden behind thick walls and layers of access control mechanisms, allowing access only to a selected crowd. The idea of open source production rather resembles an open bar camp where everybody can drop by, start working, helping others out and eventually enjoy the fruits of joint collaboration without having to get the purse out of your pockets or ask a superior for permission. Internet security production in the cases analysed resemble bar camps, ones with thick walls around them, but with an increasingly lax door policy as time passed.

In both cases, the response communities were more open and accessible than appreciated by many therein. Especially in the Conficker case, cabal members often complained that they had no control of who could access jointly produced data, who they had to collaborate with and whether these persons were trustworthy (Interview 48). But this does not imply that the Cabal was an open gathering. On the contrary, even most respected malware analysts were not granted access to the mailing list, or only after months, as they were not known by any of the Cabal members in the first place. The mailing list of the cabal rose to a few hundred members by early 2009. Access to one of the Cabal's key intermediary goods, the sinkhole data, was apparently not restricted by tight access-control mechanisms.

In the Estonian case, getting access to the inner circle of the collaborative effort required informal membership in the informal Estonian Internet security community. The IRC channel of the Estonian community and their wiki had a similar function as the mailing list in the Conficker case. The number of participants rose sharply during the incident, just as in the Conficker case. A second way to participate in the joint effort was to be a member in one of the global, most access-restricted security mailing lists. Holders of an Estonian ID card with proven Estonian language skills could also apply for access to the Estonian Security Incident Management (SIM) system.

At the process level, there are some variations. Responding to large-scale incidents requires a networked effort including many organisations, existing communities and new ad-hoc communities. A great deal of activity still happens at the organisational level, both by consumers of ICT and vendors. Patching software, updating systems with these infected systems are but two processes that are performed by organisations independent of any other networked efforts. These tasks are performed solely within an organisation, and access to this production platform requires no less than a certain role in these organisations. This also applies to the policing processes. Getting hold of the perpetrators was not a priority in either case, mitigation was at the centre of the response. But when policing activities happened, they were unsurprisingly dominated by law enforcement agencies, traditional security institutions that don't aspire to openness. Response communities have refrained from offensive vigilantism.

6.2.2 Community access

The response efforts partly relied on existing security communities. Consequentially, access to the response depended on access to these communities. By and large, the characterisation of the response efforts given above — walled bar-camps with a strict door policy — applies to communities in general, too. A security professional underlines this idea:

> "Once you break into that, once you get past the perimeter of security and you break into that inner circle of trust, there's lots and lots of sharing." (Interview 74)

Mailing-list-based communities differ in their accessibility. Therefore, statements pertaining to different degrees of accessibility are not contradictory. This view assumes relative openness:

> "It's not really a tightly closed thing. There are lots of groups in this scene, it certainly isn't closed group… There are lots of firms, organisations, and first of all people who work there, who form these loose clusters. It's not one big closed (?) thing." (Interview 37)

The security community simply varies. Some mailing lists are entirely open; a community of non-corporate researchers has only low hurdles; communities of operational security staff working for network providers however are very selective and demanding regarding their membership.

FIRST hosts several communities under its umbrella, but with layered access depending on their membership type. So-called liaison members can participate in many aspects of the organisation, and no requirements exist for this membership. When it comes to the core business of FIRST, incident handling, requirements for membership are more demanding, including a documentation of an organisation's incident handling policies and existing members sponsoring the application (Interview 43).

A mere visit to the website of one of these communities suffices to convey their seclusion. There might be a few public mailing lists like Daily Dave (cp. section 5.2.1) or the general NANOG mailing list, but those communities that appear to matter restrict access to their resources, be it mailing lists, meetings, or conferences.

While access criteria vary among the communities, they tend to comprise one or more of the following: an operationally influential position, a proven track record, professionalism, being part of the web of trust of existing members, and willingness to adhere to community rules such as engaging responsiveness ("no lurkers"), commercial appropriation or dissemination of shared information.

Access is based on professional role and therefore indirectly to the employer. On the other hand, access is specifically granted to an individual and not to a role in the company. Having an influential operational role in a company is important, but not always required, let alone a sufficing criterion. A good example is NSP-Sec. The list's focus is operational mitigation of ongoing incidents. To get things done, the community depends on a membership base that actually controls and manages the traffic flow through the Internet's backbones. A requirement for aspirants to be granted membership to the group is that they have "have the capability and authority to make operational decisions…and…have security responsibility" for large Internet networks and backbone providers (Interview 23). A different constituency wouldn't be able to obtain the objectives set by the community. In addition, less restrictive requirements would lead to a larger membership base, which would naturally increase the likelihood of leaks to the outer world by a talkative member.

The capacity to act requires not only access to decisive systems, but also the technical skills to adjust the right parameters during an incident. Communities require candidates to prove that their skill set is valuable for the community. For persons holding crucial positions at larger ICT corporations, the presence of such competence and a high degree of professionalism usually is assumed (Mathew & Cheshire, 2010, p. 6). Strangers and juniors, who haven't yet left their marks on the security world, have more difficulties in convincing others of their skills (cp. section 7.3.3 on the community socialisation of a junior member).

Communities' mailing lists are used to broadcast information, for example, about ongoing attacks. In the case of one of the common DDoS attacks, the attacked party or a community member close to it would send a message to the list, detailing the techniques used, the source and target addresses and any further information list members need to act upon such information. Community members are supposed to react to this information.

In the case of a DDoS attack for example, a network operator would check whether attack traffic originated from their networks and possibly intervene, e.g., by setting filters. Each list member is expected to take care of his technical bailiwick and "do something within your span of control & influence to fight badness" (Greene 2012a, p. 20). List members are not supposed to be mere spectators. "Lurking", as it is called, leads to bans on some lists, e.g., on NSP-Sec or OPSEC (2012a, p. 20). The TF-CSIRT community has a more relaxed approach, unsurprisingly. Its loose internal organisation was the consequence of disagreements over the EURO-CERT, when European countries could not agree on a single approach to CERT-based Internet security measures. Those communities with the "expectation to act" (2012a, p. 20) shift some substantial elements of responsibility and authority to inter-member relations. A request for help by member A from company X can in essence lead to member B from company Z, a direct competitor of company X, taking down machines in company Z's network (Interview 74). A member that does not act accordingly would fail to meet the expectations of his peers; this "inability" in turn "erodes trust and your reputation as someone who acts" (Greene 2012a, p. 20).

The security community generally has strict policies how to handle information that is shared by other members. In general, any information and data shared on a community's mailing lists must not be forwarded to non-members without consent of the person who initially originally sent the information around. These restrictive sharing policies require that a member must not share information received from other members on the list with colleagues unless explicit approval was given.

> "The correct practice on how to use the trust groups to share, actually gain the information from the trust groups within your own organization is that, when you see information being shared that would be good for your own organization, you would reach the information source and get an okay for that information to be shared." (Interview 82)

Hence, even though community membership occasionally relies on a particular role in the employing or otherwise affiliated organisation and even though a person is involved in community work in consultation with his superiors, a member must not disclose information from another list member without the latter's consent. This again leads to interesting loyalty formations. E.g., member A from company X sends valuable data, information, or intelligence to his fellow community members. Member B from company Z, the direct competitor of company X, knows that his friend and colleague across the floor works on a new product which would benefit a lot from the information B has just received. However, B would want to think twice and eventually not share the information with his colleague. If list members learned about such a breach of community confidentiality, B would presumably lose both reputation and membership with the community.

Once a person is proposed for membership, a vetting process is initiated, in which the aspirant is examined regarding the community's specific set of access criteria.

For some groups, usually mailing lists, the responsibility lies with a single maintainer of the list. In most communities however, candidates are vetted either by a group board or in a joint effort by all community members. Usually, communities require one or more existing members to vouch for a prospect, i.e., to say something along the lines of, 'Prospect ABC is totally trustworthy, I've worked with her several times and she is extremely reliable and competent.' Some groups ask their installed membership base whether they object to the admission of the prospect.

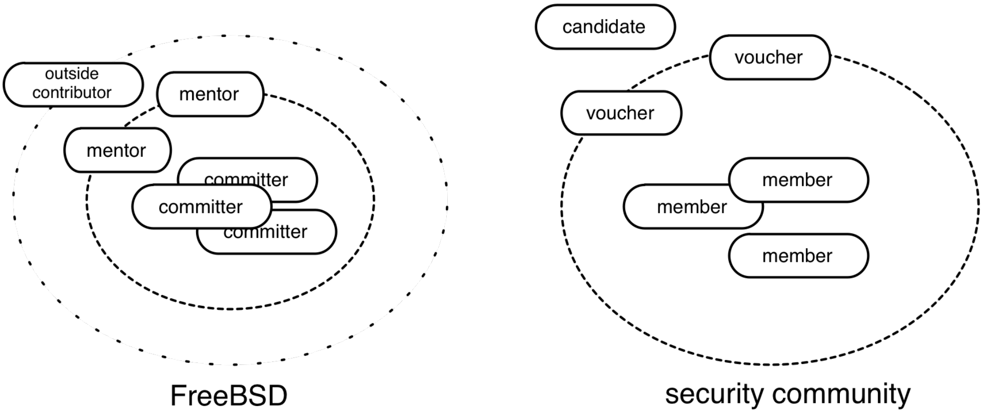

Figure 6.1: Getting access to inner-FreeBSD circles vs. an Internet security community

At the same time, prospects agree to adhere to community rules. An example from the YASML community: When the vetting referee receives negative comments about a candidate for the YASML (cp. section 5.2.3) community, the candidate gets the chance to rebut the anonymized comments, which are sent to him by the referee. The communities have obviously borrowed from academic peer review for their vetting processes. The Ops-trust community has streamlined this process with a neat online application that allow members to recommend other members and candidates. In addition, they apply a so called "meshed vetting model", in which "new members of the community to take time to document their trust of existing members — while existing members document their trust for new members" (Greene 2012a, p. 47). In some closed groups, members agree to adhere to community rules by signing NDAs or other less formal written agreements (Interview 74; Interview 43).

6.2.3 Information shared

A key characteristic of peer production is the unrestricted access to informational resources, intermediary or final results and the ability to alter and change these goods at one's own discretion. While the accessibility of processes like malware analysis, defensive DNS or situational awareness can be analysed, modifiability is a somewhat trickier analytical criterion. A process is a set of activities performed by an actor using certain resources yielding certain products. As such, a process is not copyable or forkable with a click of a mouse, just as one has seen it in OSS or information-based peer production projects. But there is a crucial element in any of these processes: informational resources used and the informational yield produced in a response effort. The borders between Informational resources on one side and informational yields on the other are somewhat blurry in large-scale incident response. As an example: One of the yields of sinkholing, lists of infected machines, is at the same time an input resource for DDoS mitigation. An analysis of openness of incident response efforts hence has to look at the informational items used, produced and shared in such a response effort.

The security community is characterised by its struggle between sharing and non-sharing, needing to share to get the job done, not wanting to share because of the risks and potential disadvantages. Information is shared with the public, within security communities or trust groups, directly with other persons in the community, or not at all.

By default, the community operates separate from the public by sharing information and knowledge within the gates of the community. With *community-only* (or higher) as the common stamp on Internet security related information, the security community hardly analyses the cost-benefit ratio of openness/secrecy for each informational item. This approach might not result in the most effective organisational form, but it reduces the risk of suffering a disruptive blow from attackers. The following paragraphs take a more detailed look at some information types that are relevant for the response against an attack. *Table 3* at the end of this section lists these information types and their respective circle of sharing. In true open source peer production, information would be shared out in the public space. The security community however defaults to share behind its community curtains. In addition, certain information is only shared on a bilateral basis or in so called high-trust groups, i.e., groups where either everyone knows everybody or any member has been and constantly is being vetted by his peer community members. Enforced sharing with national security institutions is not being discussed here.

Attack disclosure: The most basic information in incident-related Internet security questions is that of an ongoing or previous attack. Attacked organisations have opted for different disclosure paths in the past, based on their individual assessment of the consequences of disclosing an attack and details on it. If an organisation discloses an attack, the information is often only published ex-post.

> "An ISP wouldn't announce, 'We're having a mass intrusion going on here.' … In the very moment it is happening, they try to conceal the news, until it's clear what is really going on. [Until then], you wouldn't want to share it with third parties." (Interview 37)

Disclosure policies of individual organisations also depend on the jurisdiction of the respective organisation. Some countries have mandated an organisation in certain sectors to disclose attacks on certain services. In the U.S. and the Netherlands, for example, companies need to inform their customer base when the security of the latters' personal data was breached. Legislative efforts, e.g., in the European Union, likewise aim at mandatory disclosure for attack on ICT systems (European Commission 2013).

Victims' identities: Related to the disclosure of an attack are the identities of those who are negatively affected by an attack, whose systems have been brought down or intellectual property or identity stolen. This not only refers to individuals, but also to organisations who have fallen prey to a DDoS attack.

> "And then I get pings after the blogs go out saying, 'Hey, who was attacked?' … We've passed it on to law enforcement, they'll deal with it from there. There's no need for it. And I get pings from commercial companies that sell DDOS services that say, can you give me the list of the victims? It's like, absolutely not!" (Interview 85)

Disclosing the victims of attacks is regarded as further damage. In cases when attacks were successful due to negligent system administration or software development practices, disclosure can work as naming-and-shaming, increase security standards in other organisations and thereby better the overall security situation. Nevertheless, strong voices in the community suggest that victims' identities should only be disclosed to law enforcement. The general rule of secrecy-vs.-openness of the community applies here, too:

> "I think that comes back again to the idea of sharing and dissemination of information, where's the value?" (Interview 85)

Telemetry data: Among the most widely shared information is telemetry data, i.e., data originating from malware-infected systems that is gathered by sinkholes and similar traffic monitoring systems. An infected system such as a drone is both a victim of a previous attack and a likely vehicle for a future attack on other systems. Sharing such telemetry information with IP addresses of malware-infected machines usually is inevitable for responding to an ongoing attack. At the same time, telemetry data allows identification of machines that apparently have not been secured sufficiently and contain a malware which itself might be vulnerable for re-hacks. Therefore, lists of all infected machines worldwide are not shared publicly, even though it might urge some owners of large systems to clean up their infected systems. To avoid the risks involved with universal sharing, the community shares telemetry data only with those who are responsible for the infected systems or the networks that provide Internet access for those infected machines.

Some companies might choose to not end up in honeypots, telemetry data, and statistics that reveal how vulnerable their systems have been. An interviewee told me that he could not find one IP address belonging to Goldman Sachs, Charles Schwab, or some universities in Conficker honeypots. This could either be explained by a superbly managed ICT network in the financial sector (Interview 47) or by the filtering of outbound traffic from a corporate network to the allegedly predictable IP ranges of the CWG's honeypots (Interview 38).

Malware samples: Until a few years ago, AV companies handled malware samples as an asset, which they did not share with their competitors or other types of malware collectors such as researchers. Consequentially, companies built up their proprietary malware libraries. Eventually however, AV companies switched to a sharing model, arguably driven by detection rates of AV software in software reviews so low that it threatened to undermine the case for installing AV software at all. However, it can still be difficult for a third party to get access to such malware samples, especially to samples of sophisticated botnets (cp. section 7.3.3). Usually, two or three days pass by before AV companies exchange newly discovered samples with their competitors (Interview 37). The interviewee denied that delayed signature sharing, especially for botnet related malware, had a negative impact on security. While AV companies had an interest in deploying new signatures as quickly as possible, botnet researchers have a more medium term perspective, and are interested in the communicative structures of botnets. According to an interviewee, malware would usually be in the wild for weeks, if not months before it is detected (Interview 37). Only in 2012, the US government has set up a program to share malware binaries with their contractors in the military sector.

Data pertaining targeted attacks: Telemetry data originating from targeted attacks are treated differently.

> "[W]hen you are dealing with issues of a targeted activity, you could be in a position that you don't want to tell anyone in the originating country that you are aware of that activity origins." (Interview 82)

One possible interpretation is that — given all the reports on alleged state-backed China-originating intrusions — the Internet security community is biased towards Western interpretations of Internet security. This European interviewee however stressed that secrecy in such a case had no political reasons, just as the security community generally describes themselves as apolitical in their conducts of fostering Internet security: “It is not politics, really, it's security that is involved…” (Interview 82) A more matter-of-fact oriented reasoning is that details on targeted attacks increase the victim’s vulnerability. Therefore, AV companies and other community members do not share samples of targeted attacks, as these samples would reveal the identity and other characteristics of an attacked person and organisation (Interview 63).

Operational capabilities: The community applies a range of counter measures to monitor, analyse and respond to attacks. Not everything is perfect, many solutions are improvised, and therefore prone to be circumvented by the attackers or vulnerable to attacks by themselves. An interviewee replied to the question as to whether it would be possible to avoid one of their important defensive techniques:

> “Yes. We do. But again, this is not something I'd tell them, this is not something that we'd say.” (Interview 21)

The community’ sources and methods are its gem and hidden capital, upon the secrecy of which the strength of the community partly rests.

> “Actual domains, IP addresses, hashes, features of malware, what type of research people are working on. Mistakes that the bad guys make and the intelligence that we gain from that. How we're collecting malware, how we're tracking the populations of botnets ... How we're doing our work, that's the operational stuff we're talking about. Our sources and methods, and the intelligence that we gain from them screwing up. … These are the things that we would prefer the bad guys not to know.” (Interview 21)

One of the community’s contributions to ensuring the Internet’s functionality are its conducting of operations against ongoing attacks and their backers. Traditional security organisations have by default applied operational secrecy to protect their activities. And so does the community:

> “[W]e have to keep our operational details out of the hands of the adversary, so that they can't counter the ways that we fight against them” (Interview 21)

Table 6.1: Sharing of different types of information with different audiences

Type of information	Forum of sharing
Attacked systems	Some sharing with the public
Attackers' methods	Community
Attacks	Any, depends on attackee's preference
Botnet details	Community
Botnet sourcecode	Community
Data identifying victims	Community, and with law enforcement
Mailing lists	Community
Malware signatures	Any
Malware statistics	Depends on specific data: public (aggregated data); community (raw data) public (naming-and-shaming)
Malware vulnerabilities	High trust groups within the community
Operational details	Community, bilateral
Passive DNS	Community
Response teams details	Community
Stolen accounts	Partly public; community; no sharing (depending on legislation)
Targeted attacks	Private; community (not shared with originating countries)
Tools	Partly community; partly no sharing

Accordingly, the communities' general policy is that any information that tells attackers about the current state of counter measures of the response communities needs to be safeguarded.

Botnet intelligence: Given the pivotal role of botnets for any type of Internet-based insecurities, botnet intelligence is particularly sensitive. Acquiring intelligence on botnets is labour-intensive and therefore costly. Sharing such information with a group that is not thoroughly trustworthy, let alone with the wider public, can have significantly detrimental effects:

> "'I think there's a botnet here, I think it's being managed by these guys in the Ukraine, I think I can prove this, because here's the username and

> password, it's encoded in binary', and post it to the public mailing list, the bad guys will read it, go, 'Oh, look, that's our botnet', they'll move it, change the password, hide it from us…" (Interview 13)

This secretiveness on the community's behalf also comprises their tools and methods:

> "[T]he ways that we fight against them, the way that we collect their malware, the way that we lure out their botnets, the way that we take over the botnets, the fact that we're about to take over their botnets. That's absolutely gotta be obscured and protected." (Interview 21)

Such data being open would yield asymmetric benefits for the attackers. While the responding side spends three months or so[117] identifying the intricacies of a botnet, the attackers can slash that hardship by a simple update that, for example, changes the command and control structure. Consequentially, the community shields botnet intelligence away from the public. Within the community however, sharing and collaboration is vital to get botnets analysed. In the Estonian case for example, screenshots of advertisements of botnets involved in the attacks were shared among CERTs, AV companies and via communities' mailing lists (Interview 63).

Tools & services: There are numerous tools developed and shared by the security community, either publicly or within the community. Partly, these tools or services are given away as free (as in free beer) services (e.g., tools provided by security companies like Team Cymru), or as traditional open source software or service (e.g., software projects of the Honeynet Project). AV vendors however are less interested in sharing the tools that make them better at their core business, i.e., detecting and deleting malware.

> "They don't want to share their secrets." (Interview 13)

Therefore, most security companies will not publish their "new system for antivirus analysis or for fingerprinting binaries or for sandboxing or for sinkholing", while, for example, the Honeynet Project publishes source code and tools. "Data charity" (Interview 48) projects like Shadowserver publish the data they have collected, but keep their system design and implementation proprietary.

[117] Several interviewees mentioned three months as a usual time frame to analyse a sophisticated botnet. I have no indication though how much work that entails exactly. But based on the experiences with the Conficker botnet, three months of work for two persons appears to be a plausible estimate.

Non-specified types of information: Not every piece of information is withheld from the public or the wider security community to reduce the risk of use by the attackers. Some researchers choose not to share with the wider community because of direct appropriation of that information by the security or AV companies.

> "In an ideal world, I would say, I find something interesting, I would tell all the good guys. All the guys would pull together and immediately I'd do the analysis, would say, 'Right, this is what the problem is, here's how we're gonna fix it', we'd all work in the same direction. But, in reality, if I was to post it like that, an AV company would take that sample, would write a blog post about it saying, 'We have discovered this sample'" (Interview 13)

Literatures on open source theory or commons have discussed the effect of the appropriation of shared information known as free riding for quite some time. Free-riding has been interpreted as a tragedy of the commons; but open source theory has come to the conclusion that the appropriation of community products is not detrimental to contributors' motivation as the latter is mainly driven by intrinsic rewards, and the itch-to-scratch is not eaten away by a company using the community product for its commercial endeavours.

6.2.4 Openness ideology

The previous sections of this subchapter on Openness have described the state of openness in large-scale incident response and in security communities in general. Sharing might have been very common and inevitable for the response, but resources and results were hardly ever shared with the public. Only the end result of the endeavours, the re-established functionality of the Internet, was available for anyone.

The sharing practices of open source communities are mirrored by a respective ideology.[118] It is pointedly summarised in the phrase: "Open usually wins." (Andy Rubin, cited in: Stone 2010) This idea has contributed to the socialisation of computer geeks world-wide. Unsurprisingly, it has spilled over to the neighbourly domain of information security. The claim of the *Full Disclosure* mailing list — "Full-Disclosure — We believe in it" (cp. beginning of the superordinate subchapter 6.2) — is only one indication thereof. A truly open-source-like approach to information

[118] I use the term ideology non-derogatory, scientifically and neutrally defined as a system of ideas.

sharing would be to share all informational resources available and produced by members of the response communities, so that peer members can play with such data. The argument of the openness idea is that openness would eventually result in better outcomes and services — a song too familiar to students of Open Government or Open Data. However, security communities have no institutions akin to the GPL in open source communities to ensure the accessibility and commonality of their data. But openness of Internet security production not just implies access to data, but also to tools and methods required to gather, analyse, compute, or present that data. The beginnings of malware research were likewise driven by technological curiosity, less by Manichaean concepts of goodness and badness. Eventually however, this drive towards openness ended.

When it comes to openness, today's security communities favour the idea of sharing only behind the closed walls of their access-restricted mailing lists. The security communities' way of sharing deviates from open source sharing in a number of ways. In general, security communities that are relevant for incident response have an instrumental relationship to openness. Their normative goal is clearly not openness.

Second, their focus is securing the functionality of the Internet, apprehending the perpetrators and protecting victims. Responsibility towards actual or potential victims trumps openness.

> "You know, I've written blogs on DDOS attacks, and I'll clearly put out all the servers and the IP addresses that the bad guys are using for the controllers. But I won't put anything out there about the victims." (Interview 85)

This secondary role of openness because of the role of victims is also shared by groups like Shadowserver or the Honeynet Project that resemble open source projects in so many dimensions like voluntariness, non-commerciality, even building open-source software tools very helpful in incident response. These projects had their internal quarrels before they eventually converged on a stance that could be described as "responsibility first" or "defaulting to secrecy."

> "[W]e used to have some internal arguments…because some of the guys wanted to basically put everything out publicly. They said, you know, just shine a light on it all, just put it all out there. And my litmus test was, could this impact the victim? If this impacts the victim, then I don't wanna put it out there." (Interview 85)

Relatively open data projects therefore measure the eligibility of potential data recipients and adopt the traditional need-to-know principle.

> “We really take care of that only those persons get our data that are eligible.” (Interview 37)

> “The Russian CERT, which is a government entity, I give them all the data on Russia only. I'm not giving them the data of the US or Germany or anywhere else. … Russia only gets Russia, Indonesia only gets Indonesia, Malaysia only gets Malaysia.” (Interview 48)

Third, some communities seek to control what information is used for and hence only share it with recipients and projects that help or at least not stymie their cause.

> “[I]s the reason I'm putting out the data to the community because it's truly going to be valuable to someone else, in order to let them do more research and potentially better safeguard against these threats?” (Interview 85)

Such conditionality requires ex-ante filtering based on the sharers’ knowledge and judgement. The exertion of managerial control totally runs against the reasoning of the openness idea and its alleged superiority in some domains of human creativity. The basic idea behind openness and peer production is that accessibility and reusability of resources facilitates a playful use of resources in ways totally unexpected by sharers by persons with different cultural and technological mindsets and backgrounds, leading to unexpected innovations. The Conficker cabal has shared data with researchers and universities.[119] For researchers and institutions with a proven track record, getting data is straightforward, for newbies and researchers from potentially shady places there are more hindrances:

> “With other universities I am a hair more careful, so I kinda know exactly what kind of research questions they are working on.” (Interview 21)

Applying conditionality for sharing is not pervasive though. With the political sphere, the community has a relatively lax sharing modalities. The technical community does not follow a political agenda when it comes to sharing.

[119] My colleagues at Delft University of Technology are amongst them.

> "The community is sort of apolitical, in principal. It's about getting rid of a technological problem. And how the state approaches these issues, you agree or disagree with it … but in the end, it's the state who decides how to solve a technological problem. Our job just is to point at and inform about potential problems. … We can't prescribe them how to solve it. You just can't do it. Once you do that, you'd enter a big minefield. If you say, 'We only give you data, if you adhere to the following principles,' then you'll lose contact to lots of folks, as you'd really annoy them with that."" (Interview 37)

FOSS communities have debated a new GPL licence that would exclude projects and institutions involved with surveillance, military armament, or human rights abuses from the right to use such open software. In the late 1990s, the FOSS community split in two camps, with the Free Software Foundation supporting the restrictive GPL with its viral openness characteristic and the Open Software Initiative (OSI) as evangelists of free-as-in-free-beer freedom and less viral licences (Klang 2005). In that sense, security communities stand in the tradition of the OSI — minus the openness and sharing with the public at large.

But with all these limitations of openness in and by security communities, the old luring torch of openness is still smouldering.

> "You know, we need to find a way to have an open source method of sharing to help each other. Since we've been around there's been a lot more community action than has ever occurred before. … There's a dozen live working groups now that never existed before, because there's the community." (Interview 48)

But it would still be sharing behind the walls of an access-controlled network. The security community is a gated community as it subordinates openness to other objectives and values.

6.3 Socialness

While openness has been praised as the panacea to all sorts of societal ills, socialness characterises the efforts to produce the respective informational goods. The *peer* in peer production hints at the defining equal status of participants in the production endeavour. Interaction between the peers is based on social exchange rather than monetary transactions or hierarchical commands. The purpose of this subchapter is to analyse the social characteristics of the response activities in the two cases.

As described in section 2.1.4, socialness refers to the social transaction framework in which the peers exchange their contributions. The defining characteristics of socialness are the voluntary nature and intrinsic motivations of contributing actors, the irrelevance of hierarchies and markets as the driving transaction system, and a governance system that affords states and corporations a lesser role.

The subsequent section discusses the motivations of the participants. In peer production theories, the risk of appropriability of joint production has been described as a major risk to the motivation of volunteers; the section Appropriability discusses the situation in Internet security communities. Section 6.3.3 analyses the peculiar relationship between employed individuals and their affiliated organisations in response groups and security communities. Finally, section 6.3.4 discusses the relation of the response communities with the non-social transaction frameworks of markets and states.

6.3.1 Motivations

One of the puzzles of social production has been that there are persons voluntarily working on informational goods in the first place. The literature on open source and peer production has identified several intrinsic motivations that make contributors contribute despite the lack of pecuniary lures or hierarchical whips. The motivations of members of security communities are strikingly similar to those of persons in open source projects. The motivations uttered by interviewees fills a variety of intrinsic motivations known from open source projects such as the pleasure of creating new things, supporting a normative cause, and "indirect appropriation" (Benkler 2006, p. 424).[120]

Pleasure of creating. Among the several forms of intrinsic motivation, the pleasure of creating new things, fiddling with technology, and solving technological puzzles are among the more salient. Contributors participate in Internet security communities because of emotional rewards. A participant of the Conficker response group:

> "But by the same token, we had a tremendous amount of fun. There was a lot of good camaraderie amongst those people involved, working

[120] For discussions on motivation in peer production in theoretical literature, cp. section 2.1.3.

> together late nights, sharing findings, updating each other on different things that we saw, not just with the sinkholing and the statistics, but also just in looking at the malware and just in talking amongst ourselves about different things. The friendships that evolved and developed was great." (Interview 85)

Interest in technological puzzles is another driver. Hacking in general is the art of mastering technology that allows a user or a hacker to use technological artefact in ways not intended by producers. An interviewee pointed out that some community members were attracted by any forms of hacking including black- and grey-hat forms in their teenage years. With coming of age though, and increasing awareness of potential legal issues, they turned to the security community as a more appropriate environment to live out their "insatiable curiosity" (Interview 48) for technological puzzles (Interview 37). In that sense, the Internet security community poses as an alternative within the boundaries of legality to legally ambiguous grey-hat, let alone outright illegal black-hat hacking. On a more mundane practical level, community members want to "help establish sharing as a solution for security" (Interview 48), or showcase techniques to overcome the botnet problem (Interview 72).

Unique to security communities as compared to, say, open software or collaborative content-production communities is an archaic intrinsic reward that usually has no place in superficially pacified, über-organised, gender-mainstreamed, and verbosely deliberative societies. Interpol's director for cybercrime and cybersecurity, Michael Moran — appalled by the child-abuse imagery that his team needs to look at on a daily basis — went back to Ernest Hemingway to describe their motivation: "There is no hunting like hunting of men." (O'Toole 2009; Hemingway 1936) Steve Santorelli, manager at security company Team Cymru and former police officer, subscribed to that:

> "So the payoff for all of these passionate people that devote all their volunteer time is that they see bad guys going to jail." (Interview)

Another interviewee with a more technical background chimes in:

> "To have the satisfaction of knowing that at least this guy cannot get away with what he has done without actually going to jail or getting fines or anything, so that's the biggest motivation I have." (Interview 73)

The "bad guys" are not only the target of the community's operations, but also a motivator. While technical experts usually fight against technical complexities, insufficient APIs, overly ambitious customer demands, flawed documentations and flawed technological implementations,[121] the security community fights against real human antagonists.

> "It's the best opponent I've found, the bad guys. I'm intellectually curious, and I like chess!" (Interview 47)

A second pleasure-of-creating form of motivation unique to the security community is to deal with secrets.

> "If all that stuff wasn't so amazingly secret, it would be half as interesting. You just get loads of information, you shouldn't have (laughing)." (Interview 37)

Value-driven contributions. Another form of intrinsic motivation next to the pleasure-seeking mentioned above is engagement based on norms, and values. Responsibility for the Internet was mentioned by quite a number of interviewees as a driver for their participation in Internet security communities.

With their opponents often labelled as "the bad guys", the security community paints itself — at least in the early days of large-scale cybercrime — as the only actor that was actually capable of doing something about cybercrime and Internet-based attacks. These "bad guys" face the retaliation of idealistic "samaritans" — an expression used by several interviewees — that do the job no one else does.

> "And so we've grown from that attitude and basically taken that as our mission, to be that good samaritan. Because no one else is, no one else was. There are many more now, but when we started…" (Interview 48)

While traditional security institutions have pressed for greater roles and more influence in cybersecurity, in the mid-2000s, a security vacuum emerged with the rise of organised, large-scale cybercrime. With law enforcement incapable, the security community has stepped in and acts as voluntary shepherd of the endangered Internet.

> "Because if we don't do it, the Internet will break. That's the scary thing."

[121] A recent example of such problems is Apple's flawed implementation of the iCloud synchronisation technology. While in theory, everything works seamlessly, dozens of App developers had to find out the hard way that Apple's solution for synching databases is not deployable (Hamburger 2013).

Santorelli continued with a smirk, when he referred to the by-and-large voluntary nature of contributions to security communities:

> “You've got a bunch of superhero vigilantes that are keeping the wolf from the door.”
>
> “Because nobody pays this community for its work. We have day jobs and we have our community work. And we do it because we're passionate, and we do it because if we wouldn't do it, the Internet would be overrun with criminals.” (Interview)

Members of the Conficker core group mentioned similar motivations.

> “Most of us on the Conficker core group are a little bit idealists.” (Interview 86)

A proponent of non-monetary contributions explained his personal drivers:

> “It's good karma. If you build a tool to find infected computers, identify them and help the owners fix them, that's a positive thing to do. If you stop spam or if you can provide information that helps CERTs or law enforcement to shut down, say, a child pornography ring or something like that, there's a sense of personal satisfaction that you've solved ... that you're making the world a better place.” (Interview 13)

Community members not only aim at taking care of the Internet in general and the wider world, but more specifically also of their digital neighbours. A community member, who used a community-open software to analyse his network’s traffic, improved that software and its configuration files and shared the improved version with the community.

> “Well, I saved a bit of time for myself, maybe I can save somebody else a bit of time.” (Interview 43)

This motivation is well known in open source communities. Others did not want to help “somebody”, but friends. A person who contributed to the Estonian response:

> “Helping out a friend. That's the only motivation.” (Interview 73)

In distinct communities members tend to know each other quite well, as outlined in an earlier section.

Indirect appropriation. The system of mutual help however is not just built on friendship and care-giving, but also on necessity and fear — the fear of becoming the next one whose systems are being attacked. This relationship resembles the classic reciprocity of network organisations (Powell 1990, pp. 304-305).

> "The whole thing is built on that you never know when you are in the position that you need urgent help. I have been in this situation and then I know that [person A will] aid. And next time somebody calls me, I will help." (Interview 87)

In the absence of an all-securing Leviathan, actors need to help themselves by establishing a system of mutual help. This sort of help already has some of the non-altruistic motivations of a third type of motivations, not solely based on altruistic interests, even though it results in a situation that is commonly favourable.

> "The community just wants to do the right thing, because they know at some point, they might be the ones that need to have the help from the community. And we do it because it's right, we don't do it for the extra money, because we don't get any extra money. If company X says, "Hey, I'm under attack, does anybody have any way of helping me?", and I see something that I can help them with, I do it. I don't say, "Transfer me this amount of money and I will give you help", that's not how this community works. I hope it never does" (Interview 43)

People in the community help each other out,

> "because they hope that they will get the favour in return." (Interview 83)

While this tit-for-tat approach is usually implicit, it can sometimes get very explicit. An interviewee pointed out to so called "favour books" that are apparently used in networks of law enforcement professionals. In such a system, an experienced, well-connected law enforcement officer would broker informal requests for support or information crossing different jurisdictions, and note for any participants his or her contributed vs. requested assistance. As national, let-alone international police collaboration has not yet met the necessities of the speed and geographical dispersion of cyber-crime, such favour books help to overcome existing limitations of law enforcement.

Purely intrinsic motivation does not, at times, suffice to make actors contribute to common efforts without being rewarded monetarily or else being forced to by their superiors or the law. Indirect appropriation refers to economic advantages the contributors can potentially gain not from the produced good itself, but from their contributions to the production process (Benkler 2006, p. 424). Such gains might include increase in reputation, future consulting contracts, or an increase in one's HR capital. Those contributing, for example to the Conficker response, especially those in the core groups, saw their professional reputation increasing.

> "So if you were to say that you were directly involved…, it was a feather in your cap on the PR level." (Interview 85)

A number of the contributors both to the Estonian and the Conficker response have had countless appearances in media, conferences, workshops, and hearings. The general public's attention to the response however was more of an ex-post reward that, at least in the beginning, had an expected and highly likely outcome.

The race to journalists' microphones appears to be a known problem for ad-hoc response groups and communities. It highlights the potential differences between the communities' general, public-oriented interests, and the interests of companies that contribute to the response projects indirectly via their employed community members.

> "Any company is interested in being the first, in breaking news. Because it creates so much publicity." (Interview 37)

The incentives for media exposure have become particularly obvious during the Stuxnet coverage:

> "Stuxnet, for example. They had Vanity Fair, the New York Times were writing articles about Stuxnet. People wanna hear the people who investigated that botnet talk about it. A conference as RSA. So there's a massive commercial interest to be the first to market. If you find this new attack, this new technique, there's a PR machine: blog posts, conference speaking, white papers, government briefings, private briefings for executives and senior managers and for presidential advisory groups, that kind of stuff." (Interview 13)

Others were driven to improve and showcase their skills to the wider security community.

> "I wanted to build a reputation as a good hacker." (Interview 72)

Some contributors to the Estonian and Conficker responses have had a career-boost after the Conficker and Estonia response. Contributing to large-scale incident responses also creates opportunities to develop new technologies and procedures and thereby create new business opportunities. Data collection was one of the main tasks of the Conficker community, which has led to several partly commercial, partly free data collection and analysis organisations. Another form of indirect appropriation is that acquired skills, experience, and information are used to help improve the products of contributors' employers.

> "And I think you've got a lot of those kind of people in the community, where they have other related products that pay the bills. And this is an addition to that product because it gives him intelligence, information, and builds communities that they need, that may get them jobs." (Interview 86)

Summing up this section, the response to the Estonian and the Conficker incident have been driven by the very same interests and motivations as contributors of peer production projects. An itch to scratch, technological puzzles, learning new tricks, and just working together with like-minded persons on benevolent projects for idealistic causes — all these motivations that are typical of open source projects are present among security community members. A slight specialisation of the joy-to-create-type of motivation is the fascination with the communities' secrecy and the hunting of the bad guys. The self-defence and mutual-assistance argument are the only substantial deviation from the peer-production set of motivations; but they only reinforce the impression of the high degree of socialness of the response endeavours.

6.3.2 Appropriability

The literature on common goods in general and peer production of informational goods in particular has thought substantially on the problem of the potential appropriation of a common good by a single actor or a subset of the larger community. A maverick fisherman could ignore the fishing quotas set up in his village by the fishermen's community and just grab as much of the best stock as he can, to hell with the impact of fish stock in the next season. Likewise, a greedy community organizer could try to transfer common code into a proprietary resource. Restrictive copyleft licences are an effective barrier against such easy appropriation in the world of open software. The absence of risks of appropriation is seen as a prerequisite for open source projects. If contributors had to assume that all their voluntary work would end up as a means for a greedy individual to fill his pockets, motivations to contribute would likely stall in no time. But what about the appropriability of the efforts of security communities?

In a broad sense, production of Internet security in large-scale Internet security incidents means to reestablish the functionality of the Internet (cp. section 3.2.1). The result of the response activities on this level is an intangible good with commons-characteristics that cannot be appropriated. Appropriation would mean to take ownership of the collaboratively produced goods, turn them into proprietary goods and exclude others from using them. On the level of processes such as sinkholing, situational awareness, or defensive DNS exclusive ownership of jointly produced goods is not achievable either. A process is an interplay of actors, resources, rules and results that can almost by definition not be appropriated as long as actors are independent. On the level of informational resources used and results created by collaborative efforts, there are some indications of attempts of appropriation that are discussed below in this section. Regarding the communities and expert networks themselves, it certainly is conceivable that they and their activities and processes can be changed to serve certain ends more than others — just like networks can be dominated by some actors (Schmidt 2014). However, "owning" a community and its processes, if possible at all, is something entirely different that the concept of ownership of a product. What is similar though is that both types of owning allow for the exclusion of other actors and the ignoring of their interests and ends.

In his theory on peer production, Benkler argued that a commercial use of the peer-produced informational goods would not demotivate contributors as they followed intrinsic motivation. If a programmer participates in an open source project for the sheer joy of, say, hacking a segment of the Linux kernel, this joy does not go away simply because IBM uses this code for free on its System z mainframe servers, which it sells for around $1m (Robb 2010). In security communities, opinions differ. Members of data aggregation projects feel uneasy about their data being reused by commercial vendors for free:

> "Now, if you're a company that just takes our data and sells them one-to-one, sure, then indeed we resent that." (Interview 37)

> "People would then take that data, put it into their product and sell it for money. That was disappointing, because we never saw anything go back to us on that... That sort of thing makes people reluctant to share, they want to see that it's done for the right reasons." (Interview 85)

The feasibility of repacking community-produced goods and selling them to paying clients is enabled by some fundamental differences between security communities and open source projects. Most important are the lack of openness-enforcing institutions like the GPL or other copyleft licences combined with the closure of the security community. The restriction of access to security communities has the effect that many parties, who could take advantage of data produced by such security data projects, have no direct access to it and therefore rely on commercial vendors. The lack of copyleft for such data allows vendors to not place their data in publicly accessibly data repositories. The size of mailing list-based security communities is unknown, but the response communities appear to be not smaller than media open source communities (David & Shapiro, 2008, pp. 387-389). However, the closure of security communities has the effect that any appropriator needs to come from the pool of community members. A co-producer as appropriator might be psychologically harder to take than a large corporation without a face giving away your contributions without remuneration.

This potential source of appropriation and killer of motivation is governed away. Security communities have adopted the norm to not distribute community-shared information as a product (Interview 73). If a company that had "people in the community" would use community-shared data commercially

> "and get caught, that would really raise some hell and they'd get kicked out from that community and possibly from some other communities as well." (Interview 73)

> "Sometimes it did occur, where someone would utilize some piece of information or some analysis, publish it and then give it to the community. So that way they'd have a competitive edge. And when that occurred, they'd get smacked pretty hard." (Interview 48)

Security communities have established rules whether and how to use data received by sharing within the community for commercial purposes. The basic precondition is to get permission from the person who provided the data.

There are more opportunities to depreciate community effort than by appropriating data. One of the functions of the security communities is to coordinate the response against attacks and attackers. If someone follows the debates and coordination on community mailing lists, but then develops an independent strategy that harms the activities of others, it is considered a no-go. Such an action could possibly result in exclusion from the community or at least in substantial doubts for future community-based collaboration. Microsoft ran afoul of this rule once[122] to maximize their own corporate interest with possible effects for the motivation of community members:

> "Anybody who's got their own botnet research and their own takedown efforts, and their own law enforcement cooperation is now going to wonder, okay, so if I bring Microsoft in on this, are they gonna potentially take it away from me and drive the schedule, drive the strategy towards the best interest of their customers without paying any attention to the interests of my customers?" (Interview 38)

The institutions described above apply to the security communities in general. However, these findings are backed up by observations of persons involved in the Conficker response. In Conficker, controlling the message to the public was an important element of the response; and quite a few wanted to be among the first and strongest voices to be heard:

> "[M]ost of it was really about control and managing that message, because the value is in control of that message. Control of the way that the public interacts with whatever is being put out by the group." (Interview 47)

But, again, riding the community response and the PR game around does not equal appropriation of produced goods in its strict economic sense. It is more akin to what has been described as "indirect appropriation". In the Estonian case, there have been no indications of outright appropriation of jointly produced goods. In the Conficker case, the situation is more ambiguous:

[122] In spring 2012, a working group lead by Microsoft initiated the seizure of the, as Microsoft stated it, "Zeus botnet" (Boscovich 2012). Microsoft's stated primary goal was to not bring down the botnet entirely, but to primarily "inflict costs on cybercriminals". These goals conflicted with some of the other network partners with an interest in the permanent take down. Analyst and blogger Rik Ferguson of AV company Trend Micro blamed Microsoft for prematurely exposing identities of perpetrators, thereby severely harming due legal process and the ability to prosecute perpetrators (Ferguson 2012). Dutch ICT security company Fox-IT outrightly accused Microsoft of obstructing criminal investigations (de Natris 2012; Bijl 2012). Fox-IT labelled Microsoft's "Operation B71" as ineffective, short-sighted, and marketing-oriented and called it a blow to the established trust and effectiveness of the security community; Microsoft had used information shared in the community against the interests of some sharers of that information. Nevertheless, the Microsoft employee who was part of the affected community was not expelled (Interview 21).

> "Some of the people involved in Conficker, in the takedown or in the blocking, were commercial enterprises. So although they are in the opsec-community, they are also commercial bodies. So it's not surprising that some of them were looking at ways to monetize it. I think to expect otherwise is naive. I think some people were upset..." (Interview 86)

The Estonian response did not offer similar opportunities. The Conficker response involved a number of cutting-edge technologies and techniques that could hitherto be turned into commercial products.

6.3.3 Experts and company

A key feature of peer production is the absence of hierarchies in general and managerial command in particular. The lack of proper organisation status and legal personalisation of security communities makes this question a no-brainer at first sight. As the discussions in the section Distributiveness have shown this does not result in the absence of internal authority in security communities. This section discusses the shadows of management that shade onto employed community members from their affiliated organisations.

Given the motivations of contributors and the voluntary nature of their contributions, their relationship to their employing organisations appears to be irrelevant. Many members have stressed that they participated in the response endeavours voluntarily; three examples following:

> "On average I probably spend 40-50 hours a week at [name of an ICT company], sometimes as much 60, depending on what I need to do. And I spend 20-40 hours a week on [voluntary security community projects] in addition to my work." (Interview 48)

> "But at the time of Conficker working group I was full time employed by my [employer], but I did all the Conficker stuff part time, if you will. As a volunteer." (Interview 85)

> "I know a lot of law enforcement that will take vacation time and come to these events on their own dime, because they feel so passionately about it." (Interview 74)

Such voluntariness indicates independence from employers, their management, and therefore the absence of managerial command. At a second glance, the situation is more nuanced, though. The concept of voluntariness is used in two ways: first to indicate an independent autonomous status, second to describe extra-curricular activities, which are not part of an employee's official or informal job description but have been actively sought by an employee. Some community members are indeed volunteering in the broadest sense and act within the community entirely independent from managerial intervention or any shadow thereof. They report only to themselves. The vast majority however are affiliated to organisations that have some stakes in Internet security.

> "Some of them are paid to do this, but most of them it's only part of their job, it's not their day job. … And I think [in the non-written job description–] that's where it sits for a lot of people." (Interview 86)

To answer the question as to whether managerialness is absent from the incident responses it is necessary to look at the relationship between contributors and their employees.

Before diving into the issue, some methodological and epistemological clarifications are necessary. The level of analysis is somewhat shifting and in flux in this section — and in some of the previous ones, too. When discussing questions that touch upon analytical concepts such as socialness, motivations, or appropriability with interviewees, their elaborations usually referred to the Internet security community in general and not to the two incident responses analysed in this study. But as the impressions shared by interviewees have been created in such incident responses it is fair to conclude that the modus operandi of the response endeavours followed the principles of the security communities. No single interviewee has mentioned that the general remarks on the security communities do not apply to the response communities. Therefore, it appears to be fair to apply some social-scientific non-rigidness, not further question scrutinise these impressions and assume that *partes* were not that different than *toto*.

The relationship between employed community members and their host organisations is ambiguous and vague. On the one hand, institutional techniques in communities ensure and require a great deal of independence of an employee from her employer. Benkler has come to similar conclusions in a recent paper. In open source communities, he argues, community members have to ignore the clash of interests that comes from their different roles. At times, an employee-contributor even needs to outright act against employer interests — a governance norm that is necessary to make peer production sustainable and is also accepted by employers.[123] (Benkler 2013)

> "[C]ompanies that pay employees to participate in free and open source software projects have to relinquish a significant degree of control over the actions and contributions of their employees to the project." (2013, p. 224)

This holds true for Internet security communities, too. An employee has at times to ignore his contractual employment duties, follow only the rules of the community and act against the best interest of their companies:

> "We're expected to obey the rules of the community. So if the community says, 'Don't share this data,' you don't share that data. And if that data shows up on your network or in a report that your company publishes, you've broken the rules. This has happened, we kicked people out of the community because they broke the rules. They took information that was shared in the community marked 'don't share,' and it was published in their company." (Interview 21)

Despite the necessity for split loyalties, many contributors appear to enjoy either silent approval or outright backing of their host organisations. In short, they have managerial support.

> "Generally, you have a very forward thinking boss who's willing to fund you to come to these things because it's the right thing to do." (Interview 74)

> "At least, there must be some level of silent approval from the host organization, even with the persons participating in [name of community]. Because they are really working on sharing information. Also, the information is not classified, but it is sensitive in a way. And it's sensitive to the host organizations also. So I would say ... If my host organization would not accept me to be part of [name of community], I would not be there. Even if it's individuals doing cooperation, there is this organizational, at least silent approval behind that." (Interview 82)

This corporate backing is inevitable at least in some communities with an operational focus. As described in section 6.2.2, some of these communities only want

[123] For the aspect of trust in the employer-employee relationship, cf. section 7.3.3.

members who have professional roles that allow them act and administer technical systems that make up significant parts of the Internet. In such communities, members necessarily share information gained in their work time and respond to the requests of the community by reconfiguring systems that are owned by their employers. It is apparent that under such circumstances members and their host organisations need to be on the same page. And indeed, both employers and employemployees gain from this involvement in security communities. First, an employee can spend some time on things they actually like:

> "So, if you're a malware researcher, you need to take care of the antivirus signatures, but when you have other time, you're allowed to do research. You're allowed to respond to these e-mail discussions, you're allowed to be on these conferences. You're allowed to work with the community, because it's in everybody's interest that we have a dog in that fight" (Interview 74)

Second, the employing organisations have a motivated employee and can possibly get some extra value from the employer's extra-curricular engagement:

> "Most bosses…fill your time 110% with this, with stuff that they define, and then want you to do the other 20% in the other time that you don't have. And I think a lot of people sit in that kind of role. I think there's a few people that maybe this is kind of their day job, but I think it's not the default. For example, I'd be surprised if any of the law enforcement people out there have this down in their job description. I'd be stunned."[124] (Interview 86)

This statement — backed up in other conversations with the author — also implies that for many interviewees, community work is a part of their professional jobs, which also means that employers have opportunity costs, at least, for their employees' membership. This situation would hardly be sustainable if employers were not getting something of value in return. Hence, third, companies have actual advantages from the participation of their employees in these communities. To "have a dog in that fight" is inevitable for many companies with a role in ICT security and providing base Internet functionality.

[124] This statement is backed up by own observations. In an informal background chat at a European security conference, two LE officers from a high-tech crime unit told me in 2011 that they had flown in on a self-paid ticket on a holiday.

> "Well, this is relevant for commercial organisation, too, that, if you act as a lone wolf, you'll never have access to as much information as when the community is at your hand. But to get information from the community, you need to share information. No company would get along without such information exchange." (Interview 37)

While such managerial support appears logical given the advantages described above, this may not be apparent to every manager at first sight.

> "You've got two guys who are very geeky, maybe don't have great management social skills, but when they're talking about the same problem together, they'll really exchange and work on ideas. The challenge for a manager is: How do you know those side-channel, out-of-bound conversations are happening? What do you do about it? So, hopefully, good managers, they're on top of their staff, they understand who their staff have relationships with." (Interview 13)

Even if there are advantages of having their "dog in that fight", some managers might still feel uneasy about their team working with the security community and their particular rules. It requires managers to relinquish control to the employee and abstract community norms in exchange for some vague gains. So companies and their employees often go one step at a time into the world of the security communities:

> "Hopefully, over time, the company starts to see the value of the data sharing. Maybe at first you can't officially share anything, after a while, by going to events, being on mailing lists, talking to people, you start to see the benefit to you. Maybe you start to think, 'Well, it's worth it sharing more.'" (Interview 13)

Despite this managerial challenge, it doesn't come as a surprise that many big IT and Internet companies are "very supportive" to their employees' engagement in the security community (Interview 48). The substantial support by Microsoft for response communities and anti-botnet initiatives has been described on several occasions in this study; Cisco has substantially contributed to the costly hardware infrastructure of Shadowserver's "data charity"; security service vendors like Team Cymru, Kaspersky, or F-Secure and others regularly support events that bring together members of different security communities. While online collaboration hardly comes with any costs, conferencing — the communities' second pillar of communication and trust-building — is expensive. The commercial value, corporate benefits, and indirect appropriation of community membership are widely accepted:

> “There's always gonna be an understanding that whatever you work on and whatever you do, there's gonna have to be some benefit back to the parent company. Commercial dollars is what fuels this, so if people are going to be involved in these projects, and they wish for people to be involved in the projects, there needs to be some reverse benefit to the organization or to the person himself.” (Interview 85)

The support by big companies might eventually lead to a different composition of security communities. They might move from being composed of individuals to a mix of individuals and corporations to eventually even corporations only. Given the design of this study, It is not clear whether corporate membership is a historical sequel or just a different way to design a community and its membership base.

> “So, if you have some very trusted, vetted groups, companies are signed up as well, not just individuals. Some of those groups are much more company focused as well” (Interview 13)

> “Some of the people involved in Conficker, in the takedown or in the blocking, were commercial enterprises. So although they are in the OpSec community, they are also commercial bodies.” (Interview 86)

To sum up the discussion about the relationship between technical experts, who are members of security communities, and their employers: it is somewhat ambiguous, but the community norms require and ensure some degree of employed members. Some community members act autonomously and are independent from any managerial command. For the majority of the security communities, their work for the security community is related to their functional roles in their employing organisations. Norms set up by some communities are designed to ensure the sustainability of the community and protect it from degrees of corporate influencing that could eventually stall any motivation of community members. The solitary statement in the previous paragraph about “commercial bodies” as members of communities however appears to contradict the previous depictions of individual experts as the membership base of security communities. However, the statement needs to be interpreted in the context of the OpSec community, which was created as a response to the Conficker incident and the necessity to collaborate with non-vetted persons during the response. One of the goals behind this community is to create a pool or a network of trusted experts that covers as many of the larger or security-relevant companies and organisations as possible. Thinking of “commercial enterprises" in this context is necessary to evaluate the coverage of the community among security-relevant organisations. The basic unit of community membership though still is the individual.

How do these findings relate to the response activities in the two cases? After all, this section is based on statements that by and large referred to security communities in general, and not to the organisation of the responses to Conficker or Estonia 2007 in particular. Partly, this question has been answered in the methodological remarks at the beginning of this section. Furthermore, the responses to both cases have been driven by a mix of individual effort and corporate activities. It is necessary not to conflate the activities of CWG or the Estonian security community with the entire response. While decisive activities have their basis in these groupings, activities such as the hardening of systems or the rolling out of software updates were performed by organisations and individuals often independent from the steering core.

Insofar as the response activities relied on security communities, individuals arguably were the decisive agents. In that dimension, the responses resemble peer production projects, too. The individual members arguably had substantial autonomy in their activities over their employers. This does not mean that the responses were entirely in the social quadrant and did not also include hierarchical and market forces, as the subsequent chapter shows.

6.3.4 Markets and states

By definition, socialness is a founding characteristic of peer production. Consequentially, social production by definition refers to a mode of production in which transactions among the producers are not based on monetary incentives or hierarchical command as the dominant transactional framework. For the production of a good to be categorized as social production, the role of markets and hierarchies in the production process needs to be marginal, if not completely absent. The empirical observations of the two response endeavours suggest that social transactions dominated the response.

The socialness in the response is challenged in several ways, but neither the ambiguous relationship between community members and their employer, nor the market-based elements in the responses, nor the need for expensive external services negate the observation that both responses had strong elements of social production. Firstly, an analysis of details of the relationship between employed experts and their host organisations indicates that indirect market mechanisms were in play that influenced the behaviour of the volunteering experts. They are difficult to quantify — at least with the methodological instruments applied in this study — and result in some uncertainty regarding the quality of the role of companies in Internet security response communities.

The second strain of marketness in the responses however is apparent and was expected. AV and firewall vendors update signature files; software vendors update their software and create security fixes; security service providers provide their clients with threat intelligence updates and perform incident response measures; anti-DDoS appliance vendors enable their clients to deal with an ongoing DDoS attack; operators of outsourced corporate IT produce system updates and reconfigurations to achieve resilience against new vulnerabilities. All these activities were necessary parts of both response endeavours, though not necessarily integrated in the strategic coordination by the Conficker cabal or the Estonian response community around CERT EE. These activities happened because this is how the ICT market works, how systems are operated, how service level agreements between vendors and their clients are formulated. The response to Internet security incidents is characterised by a concurrence of commercial and voluntary provisioning of security services. This study has focused on the voluntary aspects of the response as a) the empirics of the obviously market-based response activities are both too trivial to answer — of course they are not social production — and b) the Conficker cabal, the Estonian response community and the global Internet security communities played the decisive role in the responses.

Another element of marketness is made relevant by the fact that Internet security production comes with a price tag even though it is to a large extent nothing more than rearranging bits and bytes. Other than writing source code for an OSS project though, some aspects of the incident response require big iron. For example, collecting all the Conficker telemetry has taken "about a terabyte per month" and "ungodly amounts of bandwidth" (Interview 21). With the absence of computational resources created by social sharing as in the ideal-type SETI@home project, these requirements translated into the need for "lots of money".[125] However, universities and large ICT corporations have donated either such computational resources directly or the money to purchase them. ICANN and registries have waived domain registration fees. At least one person has spent substantial registration fees out of his own pocket. Such donations allow social production projects to function when they need to acquire services on the market. Any larger open source

[125] The exact amount has remained undisclosed: "We do not like to tell the adversary how much money it cost us."(Interview 21)

project requires such donations — either by having Jimmy Wales looking at you for your donations on Wikipedia or by Google once paying the Mozilla Project for searches from Firefox browsers.

Aside from the absence of marketness, the second fundamental characteristic of social production concerns the role of the state, which is either absent or does not exert hierarchical authority. In both case studies, national bureaucracies and traditional security institutions were involved. In the Conficker case, law enforcement was involved, but quite a few interviewees indicated that state authorities mainly acted only as "lurkers" who received lots of information from the response group but hardly provided meaningful actions to the response.

> "There were few formal contacts with the US government as an institution, but a large number of connections through personal channels. Several researchers within the Conficker Working Group, without coordinating with others, communicated through their own social networks with the FBI, DHS, DoD and various intelligence agencies. Questions were asked about how law enforcement could help and whether the group could help law enforcement. Later, law enforcement agencies from a number of countries placed representatives on the Working Group lists so they could follow developments, but these agencies were unable or unwilling to formally contribute to the group (though collaboration with specific individuals may have occurred)." (Rendon Group 2011, p. 19)

In addition, the actions of CERT US have been described as a total failure (Bowden 2010). The above quoted DHS-sponsored Rendon Group report on the Conficker response was not published by the DHS, but leaked by the response community (Interview 47).

Some interviewees have interpreted this inactivity by state authorities as an indication for state incompetence in matters of Internet security. The beginning of signature sharing between the NSA and the so-called Defence Industrial Base (Sternstein 2012; Nakashima 2011) — suppliers of services relevant for the national security sector — was a rare public evidence that traditional security institutions were more deeply invested in Internet security issues than previously known (Interview 38). Hence, national authorities might have played the role of the dumb "lurker" not able to contribute meaningful help only to hide their expertise and capabilities. Given the stunning competencies of NSA, GCHQ and other Western signal intelligence agencies revealed by the Snowden leaks, these government agencies could have had a better situational knowledge about the Conficker incident and possibly the Estonian attacks than they appeared to at the time of the incidents. Whatever might have been going on behind the scenes in traditional security institutions, hardly any information was passed to the security communities during the incidents. Some in the community see it as the result of cultural differences between the community and state authorities:

> "[T]he military complex and law enforcement, they're not really used to open and free information sharing. The types of things that people who want to solve IT security problems solve with, the way they handle data, is very open. It's very sharing, it's like, this is what I know you need to get your job done. And when we hand that information over to law enforcement, we don't get a similar amount of information back." (Interview 43)

Even if the role of the state in the incident response was marginal, it might have been deliberatively so for another reason. The institutional design of the response effort was rooted in earlier policy decisions forbearing from regulation that is more direct. Governments in many liberal-democratic countries had liberalized both basic telecommunications infrastructure, the supply of software, equipment and technical standards. Most of the information services running over that infrastructure had been largely deregulated. Divergent as the private actors' interests were, they shared an interest in the technical well-being of the Internet, knowing that lax ICT security regulations would be tightened if the state of Internet security deteriorates. However, none of the interviewees stated that his or her contributions were driven by what remotely resembles 'shadows of hierarchies'. This type of social-scientific thinking has been entirely absent in the communities, which joined forces to provide a public good. The shadow of the state's hierarchy might have been spotted in management boards that back their staff's involvement with the security community; the motivation of community contributors however did not come from these shadowy sources. A more direct, yet mild form of state intervention though is research funding, which resulted in seemingly independent, private contributions to the Conficker response. The apt analyses of the Conficker malware created by SRI International were partly sponsored by grants from the US Army. The small print at the end of the SRI report states:

> "This material is based on work supported by the Army Research Office under Cyber-TA Grant No. W911NF-06-1-0316 and by the National Science Foundation Grant No. CNS-0716612." (Porras et al., 2009/2009a)

Irrespective of the minor contributions of state authorities to the Conficker response and the technical aspects of the Estonian response, the security community closely collaborates with law enforcement on investigations of Internet-based crime. The collaboration tries to blend two conflicting perspectives on law enforcement on Internet issues. Practitioners argue that legal experts and managers of international organisations follow an "unworkable theory of investigations", which contrasts with "what happens" in "this circle of trust and these informal networks that have naturally evolved" (Interview, anon.). The reality of effective police investigations and collaboration with technical communities differs from formal legal procedures in effectiveness.

> "The reality is this hybrid approach, very informal, very discreet, of law enforcement and industry working hand in hand, because we been to conferences together, we know each other, we trust each other. We've shared information and that information hasn't leaked out into the wider community." (Interview, anon.)

The hybrid approach also is a two-step approach. The first step is to actually find actionable intelligence about perpetrators of criminal activities on the Internet, first and foremost identifying perpetrators learning about their deeds. Much of this work appears to be done by private intelligence companies. The second step is that law enforcement then "replicate[s] what we've done in a way that is admissible to your courts and arrest" the perpetrators.

> "Nobody outside of the circle of trust ever knows about it." (Interview, anon.)

Law enforcement agencies rely on actionable input from the Internet industry and, at least for the time being, from the technical Internet security community. Both communities apparently operate at different "speed" levels. While the technical security community operates at the "speed of e-mail…or IRC", law enforcement investigations required "days, weeks, months, even years" (Interview 82). Both communities would be "getting better, day by day, at collaborating in a way that they are not destroying each other's work" (Interview 82). Nevertheless, only a "tiny percentage" of investigations would result in prison sentences (Interview 74) and only instances of cybercrime that pass a substantial monetary threshold are investigated in the first place (Interview 38).

In the Estonian case, the role of the state was more prominent. The Estonian response community was led by CERT EE, a department of the state-sponsored RIA, Estonia's Information System Development Center, directly reporting to the

Ministry of Communications. Furthermore, Estonia's Internet security community was partly the result of collaboration between various national and commercial entities with a stake in the national Internet election system. To some extent, however, it also was the result of voluntary collaboration among technical experts. The empirics of the response apparently have elements of organisational networks, epistemic communities and social productions. Just like inter-corporate networks indicated changing boundaries of the firm, trends in Internet security might signal a change of boundaries of the state.

6.4 Conclusion

The aim of the empirical analysis in this chapter has been to identify the role of peer production in Internet security incident response. Based on the operationalisation of peer production, the methodological approach and the case selection, this translated into a study as to whether the responses to the Estonian 2007 and the Conficker incident have been distributed, open and social. Given the observations described in the subchapters above, it is clear that in both cases, incident response did not conform to all aspects of the peer production model.

In the Estonian case, the overall incident response activity was by and large a decentralized effort, with specific tasks located at certain actors. Task sharing was negotiated or emerged ad-hoc and did not follow the ruling of any sort of central decision-making body. Some degree of centrality can be found in the role of CERT EE as an information hub. For the Conficker incident, the response was a collaborative endeavour by a wider and more diverse set of actors compared to the Estonian case. Security production first happened somewhat independently, coordinated via mailing lists. It later merged into a virtual ad-hoc organisation, complemented by some external actors. The Conficker response was far more global in nature. While the data is insufficient to exactly quantify and aggregate the geographical origin of response activities, the bulk of the response to the Estonian response arguably happened within the country.

Despite cultural and national differences in their core response community, both responses eventually linked up to operational Internet security communities. These have rather similar approaches when it comes to dealing with emerging security incidents. The strong unifying and underlying element in these cases is the Internet security community — despite the differences between individual communities in technical scope, geographic reach, access criteria and other characteristics. As to the criteria of distributiveness, openness, and socialness, both cases are similar, but with differences in the details as Table 6.2 indicates.

In terms of distributiveness, both responses blend elements of decentrality and distributiveness. In Estonia, CERT EE had central role in coordinating the responses at the level of its national constituency. It could, however, not act as a central authority imposing its decisions on other actors, neither within Estonia nor abroad. Much of the response work happened on the level of commercial or public organisations in an independent, autonomous way, without anyone informing a central coordinating unit.

Table 6.2: Elements of peer production in incident responses

	Estonia 2007	*Conficker*
Distributiveness	*Hybrid of decentrality and distributiveness*	
Topology of response networks	Central role of CERT EE; decentral network in EE; distributed global support	Central role of CWG; central data aggregation; decentral defensive DNS, research
Internal authority	Absence of central authoritativeness	
Deviation and internal sanctioning	(no instance of norm deviation known)	Indications of sanctioning being influenced by corporate weight
	Unclear equipotentiality of peers	
Socialness	*Strong elements of socialness*	
Motivations	Akin to OSS projects, a combination of intrinsic motivations (some of them specific to security), indirect appropriation	
Appropriability	Minor interest in media coverage	Media coverage served as PR instrument, vulnerability data as sales leads
	Limited by norm to not use shared data commercially and shared security interests	
Relation between experts and companies	Response teams driven by volunteering, individual agents, backed by some managerial support, corporate policies	
Role of state	State as co-enabler of communities	State seemingly as passive observer
Openness	*Gated openness, in practice and idea; club characteristics*	
Accessibility of the production platform	Access to local community based on previous professional and social interaction	Access to CWG core based on personal relations, to wider CWG on access to crucial resources
	Access based on skills, trust, access to resources	
Community access	Walled communities	
Information shared	Little sharing with public, frequent within communities; highly sensitive data only shared bi-laterally	
Accessibility/modifiability of produced goods	Security as common-good consumable by everyone, but not modifiable. Access to intermediary goods limited to customers and community peers	
Openness ideology	Idea of "responsible disclosure" prevalent; remnants of open-source nostalgia	

Neither was the Conficker Cabal or the later core of the CWG an authoritative security entity that could impose its will on others. The security provider that ensured the Internet's functionality rather was a network with a mix of central, decentral and distributed elements. There were pockets of central authoritative control, e.g., TLDs taking down potentially harmful domain names. The overall impression though is that in both cases individual, technical peers were driving the response. In the Conficker case, there has been anecdotal evidence that the equipotentiality among the members of the response team might have been undercut by weighty corporate influence; however, this should probably be seen more as a hypothesis for future empirical research than as a given fact. The role of national security considerations in the Conficker response has been surprising. In the Estonian case, the reluctance of Estonian or befriended experts to seek help from their technical peers in Russia is comprehensible given discussions about the sources of the attacks. As with most botnets, Conficker appears to have its roots in criminal aspirations. Independent of the intent of the botherders, mere telemetry data gained from bots pinging to sinkholes constitutes a valuable asset for those that are after vulnerable machines worldwide. Despite the frequently stated claim that the security community stayed away from politics, its actions appear to be shaped by threat perceptions that are common in the Western hemisphere.

The responses came closest to peer productions in the dimension of socialness. In both instances, the motivations of contributors resembled those common in open source projects, including the aim to realize idealistic interests, to follow personal interests or gain from indirect appropriation. Playing with secrets and hunting men are motivations alien to open source projects, but still intrinsic within their nature. Second, the social characteristic of the joint effort was not undermined by significant attempts to appropriate the collaborative outcomes. On the higher definitional level of incident-related Internet security as re-established functionality, appropriation is not possible anyway. On a lower level, regarding community-shared or community-produced intermediary goods like aggregated attack data, for example, sanctioned community norms akin to 'don't use data shared within the community for your products in a direct manner' help to ensure compliant behaviour. These norms did not prevent competition for media attention, however. Third, comparable to open source projects is the relation between individual project members and their employers. More fundamental visions of peer production would require total independence of individual contributors from their employers when it comes to work in production communities. While contributors have stressed the voluntary nature of their contributions, there is some alignment of interests between contributors and their employers. Contributors balance the requirements of the community with those of their affiliated organisations. Market forces were certainly not absent, but response activities occurred in the social transaction framework. Fourth, hierarchical commands, whether uttered by state authorities or ordered by superiors in an organisation, only played a minor role in the responses. In the Conficker case, state authorities acted at best as passive “lurkers”. In Estonia, public authorities played a significant role in the response.

Openness, or rather the restriction thereof, is probably easiest to analyse among the other criteria. The responses in both cases have not been open in the way open source production is open. At best, the activities happened in an environment that could be described as gated openness. Restrictions in the sharing of informational resources and outcomes were significant in both cases for a number of reasons. The first kind of barrier is that the response was, just as security communities in general, organised around access-restricted mailing lists or chat rooms. The core of the Estonian response relied on a mix of personal and professional networks rooting in an organically emerging local security community and previous task forces to secure Estonian Internet-based national elections. Around this core was a second layer that was easy to get access to for Estonians. These response communities linked up to more global Internet security communities with their usual access restrictions and vetting procedures. In a similar fashion, the Cabal and the core CWG were established by persons with either established direct relations, or indirect relationships, with established networks of trust within the security community. Access to the wider CWG was less restricted, as the collaboration of a wide range of actors controlling distinct resources became inevitable to successfully implement the response strategy. Second and third barriers to commons-based-peer-production-like openness are restrictions in sharing practices and related norms on sharing. The generally accepted community-wide rule is to share potentially sensitive data only directly and on a need-to-know basis. Only a fraction of the informational resources accessible inside the walls of the gated security communities were made available for the general public. The few public blog entries, reports, and repositories of aggregated vulnerability data were a far cry from the accessibility of mailing lists, discussion boards, and code repositories of open source projects. In the early days, Internet security projects used the same open approach that is common in open source projects. One can still hear some openness nostalgia in the community, and the abstract idea that some characteristics of OSS communities should be brought back into security communities. Nevertheless, it is the idea of “responsible disclosure” and a general defaulting to secrecy towards the public that shape the modalities of sharing informational resources. Fourth, access and the modifiability of community-based outcomes is limited to peers in the communities. In a way, the response activities have resulted in a public good — an increase of reliability, a reduction in the vulnerability of Internet components and thus a reduced risk of damages for Internet users; that is, an increase in Internet security. This increased security has all the characteristics of a public good. On a more mundane level, the security increase was achieved by a number of intermediary products as described in the empirical sections. In the Conficker case, e.g., proprietary OS and signature updates helped to harden vulnerable systems and relatively open analysis and research on the other.

With these findings, the answer to the question as to whether security production in the case of Internet incident response follows all elements of the peer production model is a clear-cut no. Community-based Internet incident response is a heady mix of openness and closure, of distributiveness, centrality, and decentrality, individual voluntariness and corporate calculation. In the coordinating, driving cores of the response endeavours, activities are clearly based on social currencies such as trust, intrinsic motivation, and idealistic impetus — and not on market incentives or hierarchical imposition. These cores follow the social production model as described in the section "Defining characteristics". The community-based incident responses in both cases utilize a new form of non-market and non-hierarchy-based production enabled by the rise of the Internet, and show a significant degree of distributiveness. But they do not adhere to the full openness of commons-based peer production, since they are only open behind the walls of their gated, distributed communities. What happened to openness? Has the need for secrecy, which appears to be the reverse of the security medal, killed openness? That is the puzzle which the following chapter seeks to unravel.

7 Limits of Openness

"And we trust them to be entirely human, meaning less than trustworthy."
Francis Urquhart[126]

The empirical analysis has shown that incident response in the two cases did not follow the rules of true peer production. Nevertheless, substantial elements of peer production could still be found. In both cases, the response was decentred, and input resources were widely shared, albeit within the confines of the response teams and the distinct security communities. The response was organised differently: The Estonian response did not follow market-based principles, it was a hybrid regarding its hierarchical nature, while the Conficker response was just the other way around. In both cases, the response efforts yielded outcomes that were partly proprietary, partly not.

This outcome — no peer production, but significant application of elements of peer production — raises the questions of why significant elements of peer production can be found in the response endeavour (research question Q3a), and of the primary causes of the non-application of the mode of peer production (Q3b).

To answer this question, this chapter draws on what has been written in the theoretical chapter on the viability of peer production. Section 2.1.5 contemplated the factors and preconditions that are required for peer production to be a viable alternative to traditional hierarchical and market forms of production.

The previous chapter identified the existing social dimensions in Internet security production and described the realities of sharing in the security community (cp. section 6.2.3). From the perspective of normative praise for openness, the degree of

[126] Urquhart is the fictional main character played by Ian Richardson in BBC's 1990s political trilogy "House of Cards". Dobbs, Davies & Seed, 1993/2004, min. 47:30.

secrecy in Internet security communities is sobering. The following section 7.2 analyses the hindrances of openness, and identifies the community's antagonist, the "bad guys", as a main driver towards secrecy. The section draws on the model of factors supporting open or secretive approaches developed in section 7.1.

With the antagonists identified as the main factor for secrecy, this chapter continues with a discussion of the role of trust in the security community as a mechanism that allows distributed collaboration despite the need for secrecy.

The last section of this chapter finally aims at offering explanations for the absence and application of certain elements of peer production.

7.1 Explaining openness and secrecy

In the literature review, different perspectives on secrecy and openness have been presented earlier in the theoretical chapter in section 2.3 (cp. also Swire 2004, 2005). As a reminder: whether an actor, an individual or a group, chooses openness or secrecy as the preferred operational mode, depends on a number of factors. The basic utilitarian calculation of an actor is a comparison of the costs and gains of openness versus the costs and gains of secrecy. These costs and gains are again influenced by various factors, such as the locus of expertise and authority, and the attackers' abilities to learn from failures. Actors are in addition influenced by their personal preferences based on values and norms, and by legislation and regulations.

Locus of expertise: The expertise that is required to adequately respond to Internet security incidents usually is distributed. Most obvious is the distributed quality of forensic data; that is, any data that is linked to an Internet-based attack and allows the defensive side to gain clues about the attacks, such as the techniques involved, its geographic origins, or even the identities of the perpetrators. The locus of expertise determines the distributiveness of the response approach, not necessarily the degree of openness. As costs of organisation increase with the distribution of production, distributed locus of expertise probably favours an open production and governance approach.

Locus of systems authority: As the response to security incidents eventually requires changes in the configuration of technical systems, the locus of systems authority is a decisive factor for the openness question. Control over technical systems is extremely distributed, residing with the owner of individual technical components. Coordinating response to attacks among thousands of potentially relevant systems owners might be difficult in a closed response approach.

Gains of disclosure: Irrespective of its costs, the security community would surely gain from a disclosure of its methods, technologies, tools, and its inner workings. It is safe to assume that hundreds if not thousands of computer scientists and security experts, who were not affiliated with the response communities, would have become interested in contributing their expertise to the response teams if they knew about their efforts. Disclosing attack information would incentivize affected organisations to increase their security standards. Disclosed technical and organisational architecture would give like-minded contributors the chance to increase standards and suggest improved versions. However, these gains might be equalled by potentially even higher costs. The calculation of gains changes with different levels of analysis. Some argue that more openness could harm certain future victims, but that in the long run the security situation would be improved. Others refute this argument as entirely based on assumptions, suggesting that the gains of openness are hypothetical while the costs would be real.

Effectiveness of initial attack: Whether an attack on an information system becomes a success for the attacker depends on his acquired knowledge about the attacked system, its design, and its security mechanisms. Stuxnet showcased how thorough intelligence gathering, online and offline, facilitates a successful initial attack. The effectiveness of the initial attack depends both on information that the target reveals about itself and that the attacker manages to acquire without knowledge, let alone the consent of the attacked party. The ability to acquire intelligence depends on the attacker's intelligence capabilities. Therefore, the more the field of Internet security is filled with powerful actors like states or cybercrime gangs, the more likely is the effectiveness of the initial attack.

Attackers' ability to learn: Another argument against openness is the attacker's ability to learn from the defensive side, and their response strategies and techniques. With the defensive side publicising all their intelligence about the attackers and their own defence techniques and technologies, attackers can find numerous clues as to how to adapt their attack techniques and strategies to circumvent the defenders' stronghold and go instead for their 'soft underbelly'.

Competitive advantages: Actors can prefer secrecy also for a number of competitive reasons. Applying his theory of disclosure, Peter Swire found a number of incentives for owners and vendors of software to not publish their software source codes, the configuration they use, and the overall architectural design and the specific role of open source systems therein (2005, p. 1335). The interest of the AV and ICT security industry in sharing with other members of the security community is, at least theoretically, limited. The economic need for a workable value chain needs them to have some "secret sauce" (cp. 2.3.1) over whatever open or non-open information.

Table 7.1: Impact of factors facilitating secrecy/openness on security community

Locus of expertise	Attacks are highly distributed; although there are some knowledge clusters
Locus of systems authority	Authority over technical systems highly distributed
Gains of disclosure	Disclosure of methods, technologies, tools, humans, communication might bring some benefits, but risky and costly
Effectiveness of first attack	Mediocre in both cases (though: potentially extremely high for targeted attacks, e.g., Stuxnet)
Attackers' ability to learn	Both Estonia and Conficker as cat-and-mouse game between attackers and defenders
Competitive advantage	"Solidify positions"
Regulative requirements	Competition/privacy/anti-trust/corporate regulation as incentives for secretive sharing
Norms	"Hold your cards close to your chest", "responsible disclosure", "info needs to be free", openness or secrecy ideology (pirate party vs. national intelligence community), "the bad guys", LE mentality; institutional path

Regulative requirements: In most countries, handling of data is regulated to some extent. Public regulation and laws usually limit sharing of data on individuals, including data traces that users leave behind when using the Internet. Furthermore, anti-trust law has traditionally depreciated sharing of relevant information among companies in the same sector if such sharing would reduce competition. In addition to that, corporate law requires listed companies to share information pertaining their market performance indiscriminately.

Norms: Actors' stance on openness and secrecy as a guideline for operational design is also influenced by their overall values and norms. Visions here range from the libertarian "information needs to be free" idea widely shared among Open x activists or Pirate Party members, and its equivalent in the security world, the "naming-and-shaming" principle; to the "responsible disclosure" norm predominant in the information security community; to the "need-to-know" principle that dominates traditional security organisations.

7.2 Drivers of secrecy

As the previous section has shown, most types of information are not shared with the wider public. Most information is only shared within the community, some is only shared in high-trust groups, in a bilateral manner, or not at all. But what is the main driver for the community's secrecy? Do they exclude the public at large because governments force them to do so? Is it for competitive reasons that community members do not share? Or is it the ominous "bad guys" that make the community opt for secrecy?[127]

The answer is decisive for the viability of open modes of security production. The behaviour and demands of national security bureaucracies can, at least in theory, be influenced by regulators. Competitors can be forced by regulators to share certain information with their competitors. As long as the reasons why a responding actor opts for secrecy — e.g., distribution of gains of openness — are roughly within the regulatory reach of the community or any other national body, these secrecy-drivers can be moderated. If however the drivers for secrecy are beyond regulatory reach, changing the secrecy culture is likely to become harder, if possible at all.

The security community's secrecy is foremost driven by the existence of their antagonists. The "bad guys" do not just cause the community to respond to incidents in the first place; they also shape the community's behaviour, their norms, and practices. This does not imply that there are no other drivers that cause actors to prefer not to exchange information. Sometimes information has been held back for selfish economic interests; at others it has not been exchanged because of longstanding community norms. Overall, however, the shadow of the "bad guys" has been the main reason for the security community's avoidance of the limelight, and preference for the walled-in space of its access restricted mailing lists.

The following subsections aim to discuss these various facilitating factors of openness and secrecy.

7.2.1 States

The initial assumption for a case relevant to national security was that national authorities and traditional security organisations would likely play an enhanced role

[127] Secrecy here means that the production platform and communicational space is not open and accessible for anyone; that informational goods required for the production are not accessible for anyone; and that some produced intermediary informational goods are not accessible for anyone.

in the response to an attack. However, the Estonian response was the product of the collaboration of the Estonian Internet security community with their international counterparts. In general, states can influence the organisation of Internet security provisioning and the predominance of secrecy in several ways. State-owned organisations can participate in the incident response and push through their secrecy norms in collaboration with other actors (Mueller et al., 2013). Alternatively, secrecy can be the outcome of state-sponsored regulation and norms. All in all, states' policies and institutions have apparently not led to a significant application of secrecy or openness in the two cases.

In the Conficker case, LEAs pressured for cautious data sharing by the CWG.

> "[L]aw enforcement, once they kinda got wind of that [the data transfer to the Chinese; AS], they kind of created a big problem, because they had the understanding that we were tightly controlling the data.... And therefore, when word of that got out, that created a bit of an issue with law enforcement. And not just US, but also international. That was kind of a key issue." (Interview 85)

This cautious approach to data sharing with other countries was already accepted by the CWG's core group. Neither US nor international law enforcement substantially contributed to the response effort. Nevertheless, they wanted to make sure that the CWG had a grip on its vast data sets. Accordingly, the rupture within the CWG over the question by whom, how, and which data should be transferred to the foreign countries and the Chinese in particular led to concerns within US LEAs and raised the question whether the CWG had a leak (Interview 85).

It is not clear whether international political considerations played a role within the security communities that took care of the Estonian incident. Only late in my interview series did I realize that none of my interviewees had mentioned any sort of collaboration with Russian technical experts. For a technical community that claims to be apolitical when it comes to the security of the Internet, one would expect this community to reach out to their Russian peers to get further insight into what was happening in the attacks. An Eastern European expert said that it was possible that Russian experts were excluded on purpose (Interview 82). Western-Russian Internet security community relations are still an unwritten history.

The community apparently wishes to stay clear from state influence and not be directed by state authorities. Government involvement is seen as a gateway for governments to exert control over the community and to change its focus point from international to national security interests.

> “[A] lot of the members of the CWG would prefer it to have an international Internet interest as opposed to a national interest.” (Interview 21)

For the community it is crucial to be loyal to the global Internet as a whole, rather than to aspects of it that serve distinct national interests.

There are, in summary, only minor indications that states have significantly altered the way the security communities responded in the cases analysed. Even in the Estonian case, in which apparently national security interests were endangered, the military hardly played a role and did not affect the organisation and procedures of the response endeavour. Beyond the concrete cases, interviewees raised a number of issues in the apparently complicated relationship between states and these communities.

While the community is not keen on embracing governments, governments are unsure how to best deal with the community. Governments are used to cooperating with private companies as providers of essential services within their national borders. However, depending on the services of a transnational, trust-based network, that is not even a pure industry network, is something which governments do not often have to deal with.

> “Dealing with people on other parts of the planet, dealing with jurisdictions, it makes things very complex. It's not traditional governance in what most government people think, sort of top down, we have our nicely drawn borders around the country…” (Interview 86)

The fundamental difference between traditional security organisations and the Internet security community is that the organising principle of the former is hierarchical authority, while the latter is foremost characterised by the trust-based relations among its members.

> “It's in many ways just personal relationships. That's basically the way it works, and governments can't handle that.” (Interview 15)

In addition, communities lack the hierarchies that top-level policy makers might expect when dealing with public service providers. There is no CEO, General, or Head of Security Community. It does not even have spokespersons. It is a diverse network of communities, which policy makers apparently ignore. Instead, they resort to the traditional public-private partnership model.

> "'We need to have more public-private partnerships, we need to have more information sharing!' — I have heard that mantra be said for the last 8, 10 years. Now, we in the IT security community have been doing it all along… But, as you know, the military complex and law enforcement, they're not really used to open and free information sharing. … And when we hand that information over to law enforcement, we don't get a similar amount of information back" (Interview 43)

Traditional security institutions have not embraced reciprocal information sharing so far. Attempts by the security community to convince policy makers to follow the community's model have also failed (Interview 43). The networked state is a fast-selling item only in political science, not in politics. However, designing a governance architecture that factors in the need for a networked approach is more difficult:

> "[Governments] understand that there is something of a paradigm shift here. They've got to figure out their way through this, because nobody is going to do it for them." (Interview 86)

National security thinking already hinders security communities from exchanging information at their own discretion. In some European countries, laws govern the way in which information can be exchanged across borders once a computer system is classified as relevant for national security matters (Interview 73). Worse, the national security perspective in combination with an offensive defence policy requires state-tolerated, if not state-created Internet insecurities.

> "[C]ountries… have announced that the best defence is a good offence, and that if they're attacked, they plan to counterattack. And if you're going to counterattack and your opponent has a botnet, then your counter attack will have to have a botnet of its own. And I'm not sure exactly whether my government knows this, but it's the logical conclusion." (Interview 38)

Conceptualised in this way, national Internet security requires Internet insecurities, preferably on your opponent's side.

An example of how state regulation can unintentionally nurture the community's secrecy are due process rules. Response communities try to avoid the limelight, thereby reducing their risk of exposure to the criminal attacking gangs. But ironically, those in the security community who regularly cooperate with law enforcement and provide them with clues and evidence for cybercrime cases see their operational security endangered by some elements of due process, namely disclosure law. In most countries with due process of law, substantial amounts of the information that could serve as evidence against or in favour of the defendant needs to be disclosed to the latter. Consequentially, security researchers apparently withhold certain information from law enforcement to protect their operational security. In this case, the effects of disclosure law are comparable to being infiltrated by the "bad guys". It gives defendant attackers the chance to learn about the methods, techniques, and organisation of some parts of the security community.

> "Working with a cop… on one of these cases, in theory, he has to basically print out every e-mail that he's had with you, every Jabber conversation, every IRC/ICQ conversation, his notes from his phone calls, everything. … And maybe, if you have something that is particularly sensitive, you shouldn't tell it to the cops. Because it would put them in a very difficult position and they would rather not know. So it's not just a trust for infiltration, it's a trust in the way intelligence and evidence and information generally is handled by the parties, 'cause they're all subject to different regulation and legislation." (Interview, anon.)

The community has a history of prioritising the need for results over adhering to formalities and official procedures. The nucleus of the Estonian community was built by bank IT security experts, who joined forces even though Estonian law initially outlawed such operational collaboration. In the Netherlands, IP addresses were regarded as personal information, and the security community was thus not allowed to exchange mere log files until a few years ago. Today, as mentioned above, regulation in some countries forbids the exchange of telemetry data that includes information about ICT systems with a 'critical infrastructure' or 'national security' seal on them. Experienced interviewees, especially those with a public role for the community or their affiliated companies, stressed that information exchanges and other activities stayed entirely within legal constraints. One with a background role however replied to the question whether laws would hinder information exchange:

> "In some cases. Usually, there are ways to work around the law." (Interview 73)

Furthermore, communities would in general still operate the way they did before the states entered the ring of information security.

7.2.2 Business

Securing and gaining competitive advantages is one of the core drivers for private actors, individuals and companies alike to keep information close to their chests and resort to secrecy. Commercial interests surely played a role in how actors behaved during the Estonia incident and even more so in the Conficker response. The effects of commercial interests regarding secrecy, i.e., sharing of input and output resources and access to the production platform, remain arguable. One could hypothesise that the antagonists are to blame for secrecy more than the commercial interests are (see the subsequent section). As a second hypothesis however, commercial interests directly lead to reduced sharing and greater suspicion within the community.

Secrecy can be used as a tool to influence the relational power among the actors involved. The installation of restricted sub-groups in the Conficker Working Group institutionalised and enhanced the core group's previous status within the CWG. Hereafter, its members were the only ones who could get a bigger picture as only they had access to all the subgroups.

> "That's how [secrecy was] used. …Not so much to create, but to enforce or solidify positions using secrets or creating secrets." (Interview 47)

Whatever the core's elaborated position was actually used for, those who were in it had superior situational intelligence compared to those who were silo-ed in one of the subgroups. Moreover, they were in a better position to influence the behaviour of the remaining CWG. While its compartmentalisation allowed the CWG to scale up and reduce the risks of sharing to a lower level, it allowed CWG's core group to reduce the access to resources for previously influential members.

Just as secrecy is an instrument to "solidify positions", openness is a means to undermine the advantages of the privileged ones. An actor's stance regarding openness is influenced by her interests, normative preferences, and given access to resources. Calls for openness hence are not necessarily an indicator that a supporter of openness deems openness as a societally preferable mode of interaction.

> "I think people that are in it, you know, those that don't have the data want to see more sharing and more openness. And I think those that do have the data wanna see it more controlled." (Interview 85)

Whether one calls for openness, depends on factors such as position, interests and control of resources. AV vendors use the community to increase their knowledge, use that knowledge for their core business and increase their value for customers by adding their distinctive "secret sauce" onto that shared common.

> "[T]hey don't want to share their secrets. … [W]e publish them, we tell people what we're doing. Other groups don't, because they want to keep an edge, they want to keep it private." (Interview 13)

Quite a few in the security community however argue that the non-sharing of methods and raw-data results in non-optimal degrees of sharing, imperfect informational awareness, and consequentially a non-optimal degree of Internet security. It is apparent, however, that commercial interests and openness have an uneasy relation in the information security domain. Maintaining openness is easier if you're not trying to make money from your activities.

> "So there's no financial incentive not to be open. … I don't sell a product; I don't have any money that I'm trying to protect. I don't have competitors who are trying to undermine my market share. It's not somebody trying to get my next product features from me, so they can build them first and beat me to market. … So it's easy to be open when you're not having to hide. I'm not trying to sell the government a ten million dollar contract for secret spying on civilians or something…" (Interview 13)

To answer the question of whether non-openness leads to less Internet security could be the research question of another PhD project. The sentiment is shared by a few in the community, though.

> "[Security companies] are overly secretive, and it's purely for the concept of commercialization. Right? And that's part of why we fit in there as well ... I consider one of the largest issues the commercialization of the security field and the data." (Interview 48)

7.2.3 Norms

The security community is driven by some conflicting norms pertaining to secrecy and openness. On the one hand, many appear to be driven by both traditional and security-specific Internet ethics and policies like: the Internet needs to remain a free and open place, enabling uncensored freedom of expression; the Internet security should not be a tool of the content industry; the security community should remain politically independent, serving the technical integrity of the global Internet rather than a single company or country; cooperation with law enforcement is important to bring down attacks on Internet systems; and getting funded from traditional security organisation like the DHS, DoD, DARPA, is in line with the community's values, as it helps to get the job done. There appear to be some potential conflicts between these values, not only within the community, but also for unique actors therein.

> “I think the goal of the security community is to be more open, but I think that's more of a philosophy than an actual practice. The security community says that they wanna be more open, but I think in practice they are not as open as maybe everyone would like to see.” (Interview 85)

Later in the interview, he said:

> “[To hold your cards close to your chest] is kind of a foundational philosophy of the Infosec community. Information needs to go out, but, you know, what information to who and how?” (Interview 85)

Debates on Open Data are an example of both the social contingency and persistency of norms. State bureaucracies have for centuries largely operated as black boxes from the perspective of citizens. Their internal workings have not been exposed to the wider public. Proponents of Open Data and other Open X ideas claim that openness in public policy should be the default approach. The security community begs to differ as the following section shows.

7.2.4 Antagonists

In both the Estonian and even more so in the Conficker case, the attackers have carefully watched the publicly apparent actions of the response side and readjusted their attack strategies accordingly. In both cases, participants in the response endeavour felt they were being observed by the attackers. They had seen strong indications that the attackers were adapting their attack techniques to the actions of the defence side. Consequently, attackers would have been helped by publicly available information about the responders, their capability, and their actions. However, the risks of openness does not only comprise the attackers’ ability to learn and improve their attack techniques, and the undermining of the response teams. It stretches to the very health and personal integrity of those involved in cybercrime investigations and incident response.

The responding actors didn’t resort to secrecy as a consequence of attackers’ use of openly available information. In both cases, the responders rather applied secrecy as it is common for such Internet security and incident response communities. Secrecy towards non-members is the community's default policy.

> "[To hold your cards close to your chest] is kind of a foundational philosophy of the Infosec community. Information needs to go out, but, you know, what information to who and how? That's very, very important." (Interview 85)

The community deems the risks of openness higher than its potential merits. The more open the community acts, the more opportunities there are for an attacker to identify any technical and organisational vulnerabilities within the security community. Such knowledge fundamentally increases the attackers' chance to succeed in their attacks. Furthermore, attackers can learn how to avoid the actions of the response side:

> '[T]he bad guys watch our communications, so if I publish on an open mailing list, "I think there's a botnet here, I think it's been managed by these guys in the Ukraine, I think I can prove this, because here's the username and password, it's encoded in binary", and post it to the public mailing list, the bad guys will read it, go, "Oh, look, that's our botnet", they'll move it, change the password, hide it from us.' (Interview 13)

Some even go so far as to assume that the security community can get undermined and intruded by the "bad guys". Not everyone in the community shares this threat perception, though. Nevertheless, a trend in the security community that started with the Conficker Working Group is to compartmentalise the virtual response organisation and split response teams up into several sub-groups.

> "Initially, because we felt that Conficker was so much of a threat, we said, well, the more eyes and ears, the more people with their feet in the trenches working on it, the better it is. So that's why we kind of opened it up .... But it soon grew to, I would say, over a hundred people. And that was when it soon lost control. Because you need to classify some of the data and some of the things we were working on. Because we came to the very real distinct possibility, hey, there might even be the Conficker authors that are part of this group." (Interview 85)

From that perspective, the moment of relative openness in the beginning of the Conficker response can be interpreted as the least bureaucratic approach in a moment of overstretching, and the need for more response resources.[128]

Next to nurturing the attackers' attack skills and increasing the risk of being undermined, openness also arguably increases the risks for the persons involved in Internet security. These risks range from reduced privacy to consequences for their personal health.

> "If you're an analyst in Russia, analysing Russian malware, you might not exactly want to piss the other side off by going public, 'This and that malware is of Russian origin and done by these guys'. In Russia, you might get killed for that." (Interview 73)

This feeling of being threatened by the Russian underground economy appears to be shared by the security community, though in different degrees. An Eastern European interviewee added with a smirk that the fear of the Russian is reciprocal to the distance from their border:

> "It becomes less, if you are closer to the Russian border. If you are far from the Russian border, it goes higher and higher. So, persons from the US, they are pretty afraid of." (Interview 87)

Fears of retaliation by the Russian underground economy had been nurtured by an incident in the Ukraine. The hosting company UA Hosting had a decent share of its customer base in the criminal sector. With a slow and unguarded abuse-department, it built up a reputation among its peer hosting companies as friendly to criminals. Facing the risk of de-peering, UA Hosting decided to take down the websites of their criminal customers. The next day, their computer centre was set on fire (Interview 87).

Despite such fears, the AV company Kaspersky has posted full names and pictures of some of their apparently Russian malware researchers on their website. A senior security expert added that the community's secrecy would just be a precaution, and a greater degree of openness would likely not result in harm.

[128] This Conficker core group member would have been right to be suspicious of the Conficker authors' participation in the Conficker response team, if the theory of Bumgarner was right that Conficker was the launch pad for Stuxnet (Finkle 2011). This theory has also been shared by a minority in the community — persons with a reputation as a sober, insightful, experienced analytics of all matters Internet security.

> "Obviously the attackers know that we know where their CNC servers are. [T]hey know we know. [T]hey might attack us more directly because they know who we are and they might want to retaliate. I don't know. It's just a precaution. There might not have been any real damage done even if we would have share more information publicly. We draw the lines somewhere." (Interview 63)

The current default to secrecy is arbitrary. But given the stakes, the security community is understandably reluctant to push openness to its limits by trial and error.

7.3 Trust in the security community

The communities are secretive in several ways. First, public information about them has been scarce. Apart from a few ad-hoc groups dedicated to a temporary incident (such as the CWG), communities as network virtual organisation are usually not present in the public discourse, while individual members, especially from the security industry, are prominent in the media. Second, community resources are, with the exception of a few open mailing lists, not accessible for non-members. Secrecy towards the public and non-community members is justified by the presence of antagonists, the high costs of appropriation or sabotage of community work by malevolent persons, and the sensitivity of shared information (privacy, regulatory requirements).

A truly open form of production is apparently hindered primarily by the very existence of the antagonists. It is they who cause the security community to organise itself behind doors closed for the public at large. In the early 2000s, a number of Internet security projects operated similarly to open source communities. Communities close to the industry like those controlling the Internet's backbones have always been an access-restricted event. With the emergence of large-scale, increasingly sophisticated Internet-based crime, however, the community has thrown up walls around themselves. Nevertheless, much of the initial "hippie-stuff"-type of motivation has apparently remained in place and been transferred to those who joined later. Within the communities, commercial interests, state regulation, and the emergence of national Internet security politics have had an ambiguous, probably mostly negative, but hard-to-measure impact on the sharing of input resources and output.

What is apparent however is the qualitative impact of the community's antagonists, the proverbial "bad guys." They require the community to apply another organising social principle that plays only a subordinate role in ideal-type peer production or open source communities: trust. The need for secrecy, nurtured by the vulnerabilities created by being open, requires higher degrees of trust among those contributing to Internet security in the communities. The need for trust entails a number of institutional features that result in a production model that borrows, but also deviates from standard peer and open source production.

No community member describes the community without using the word trust. The communities' glue is trust. A network operator cutting the Internet connection of one of his customers needs to have a great deal of trust in his community colleague, who had asked for such a measure, despite working for a competing company.

7.3.1 Theories of trust

The pivotal role trust plays in the security community requires some theoretical considerations. The conceptual model that has guided this research project did not include trust as a variable that would somehow influence the applicability of peer production. Benkler himself chose to disregard non-economic discussions of distributed collaboration, in which trust plays a significant role (Benkler 2002, p. 400).[129]

Trust is usually defined as the expectation that another person will act in an expected, benevolent way, despite the person's capability to inflict substantial harm on the trusting person. It is a "risky advance concession" (Luhmann 1968/2014, p. 27). The precondition for deep trust usually is that it requires repeated interactions that allow the truster to build confidence in her expectations of the trustee's future behaviour. Obviously, there are varying degrees of trust, and scientific literature has developed respective models of varying trust, ranging from high distrust, to low trust to high trust (Jones and George 1998, citied by Newell & Swan, 2000, p. 1297). Next to this distrust-trust continuum, the literature on trust has identified three sub-types of trust: companion trust is based on personal friendships, competence trust on the perception of another person's skills, and commitment trust on contractual agreements (2000, pp. 1295-1297).

[129] Benkler discusses Adler 2001, p. 218, which proposes trust as the third coordination mechanisms next to price and authority, and communities as the third organisational form next the markets and hierarchies.

The importance of trust for distributed collaborative networks has been highlighted in previous research. In a study analysing the effects of trust on cross-company cooperation, Schilcher et al. (2011) conclude that the level of trust defines the success of such cooperation. They observe that trust was sequentially built up in such cooperation. The usual process of building trust started with an upfront credit of trust after seeing some trustworthiness in the other; with increasing familiarity trust would develop and eventually be strengthened and lead to increasing risk tolerance.

Other authors have argued that swift trust is a prerequisite for open source communities, as it allows strangers to collaborate with one other, each assuming that they have goals in common (Osterloh & Rota, 2004). Swift trust is a mode of trust that is not based on acquaintance, but on assumed common values and goals (Adler & Heckscher, 2006; Osterloh & Rota, 2004).

7.3.2 The community's trust model

Companion and competence trust are common and required among members in the security community. Commitment trust is also present, but in less explicit ways. While community members don't sign contractual agreements that require them to deliver certain services in certain times, they implicitly or explicitly agree on the written or unwritten norms of the community. Furthermore, there is an indirect commitment that nurtures trust within the community, as members usually work for organisations or companies that feel the "shadow of hierarchy" and therefore likely feel the need to deliver and increase Internet security.

Following this trust-maturity model, the Internet security community is a high trust environment. Not only is access to core communities only granted when initial trustworthiness exists; access to core communities and groups usually requires years of previous collaboration. Cooperation among network administrators is based upon a similar concept of trust. After an initial assessment of trustworthiness, the first cooperation establishes it and results in a process of ever-deepening trust, if actors live up to each other's expectations both in terms of capabilities and motivation (Mathew & Cheshire, 2010, p. 6).

Swift trust is apparently insufficient for security communities. Trust in these groups is more deeply-rooted. This differentiates them from other forms of "collaborative communities," in which "swift trust" is entirely sufficient to facilitate global distributed collaboration among persons of various backgrounds and no history of previous interactions. In Internet security communities, trust among members is based on more dimensions and components than swift trust.

Table 7.2: Components of trust in peer production vs. Internet security communities[130]

	Peer production communities	Internet security communities
Sources	"Generalised reciprocity" assumed common values and goals	*Familiarity through repeated interaction* *Calculation based on interests* *Norms that create predictability and trustworthiness*
Mechanisms		*Direct interpersonal contact* *Reputation* *Institutional context*
Objects	*Individuals*	*Individuals* *Systems* (trust in web-of-trusts, trust groups, enforcement) *Collectivity* (trust in the respective community)
Bases	*Competence*	*Contractual trust* (community policies) *Competence* *Benevolence*, *integrity*, *openness*

A community that declares trust to be its founding principle and is populated by very smart people has developed its own ideas of the model of trust which works best for them. Many see the community as a web of trust. It is too large to know everyone, but the network of trust relationships enables anyone to use trusted intermediaries to establish ad-hoc contact with any person whose input is required to solve a certain problem.

With regard to trust, an interview highlighted the different concepts of trust applied in the community. A group controlling a small number of systems providing scarce Internet services (the root operators group) would be characterised by "Trust with a capital T", a group where "we all trust each other". This trust group of the rulers of the root relies on deep trust, built up over years of interaction and familiarity among a couple of dozen of people.

> "This is a couple of dozen people, and we all know each other, we've all drank with each other, we all have met regularly." (Interview 86)

Next to this inner circle, a group that truly controls some of the Internet's existing kill switches.

[130] Dimensions (left column) and names of components in italics taken from table "Dimensions and Components of Trust" in Adler 2001, p. 218.

The security community's issue groups differed from these Trust groups, the trust level within the group is lower. They follow a web-of-trust model.[131] Instead of a configuration where each person trusts everyone else, a network of trust is created.

> "[E]verybody on here [i.e., issue groups] is people we trust, or people who've been vouched for by people we trust, or people who've been vouched for by people who've been vouched for by people we trust." (Interview 86)

It is common not to know everyone in such issue groups, even for persons who have been part of the community for years. Nevertheless, all the direct or indirect trust relationships between the actors create what Barry Raveendran Greene calls a "sphere of trust" (Greene 2012b).

A special type of web-of-trust model is applied at the OPSEC Trust community. What they call "peer meshed" is the mutual vetting of candidates by members, members by candidates, and members by members. The goal of this approach is to ensure the high quality and trustworthiness of members.

When the community needs a high-trust group, but also hundreds of participants, it combines trust groups with webs of trust. Scaling is achieved by establishing subgroups that are managed, steered and directed by the core group. This model was established during the Conficker crisis and was lauded by many as the model for future interventions.

Next to the technique of separation of high-trust groups and webs of trust as a means to overcome the scaling problem, the community regionalises its subcommunities to create the group sizes necessary to achieve the required level of trust.

Interviewees mentioned different numbers when asked about adequate community sizes. What could be described as the core community only encompasses around 500 persons.

> "A lot of the people I will see on these close groups, I can go to any number of other close groups I'm on, and the same people will be on

[131] Barry Raveendran Greene (ISC) suggested a "chain of trust" as a model. The term aptly describes how trust is lent from one person to another. The overall network configuration however is more aptly captured by the term web-of-trust. It is a characteristic of networks nodes are usually not directly connected to any other node (Greene 2012b).

> pretty much all of them. So it's a very ... 500 is probably being very optimistic." (Interview 86)

This number certainly does not encompass any person around the world having a role in Internet security. It most likely includes liaison persons from major players in the ICT industry, network providers, the security industry, and response organisation.

Community members consider the scaling of trust as an issue for security communities. The larger a community, the less it serves its raison d'être, which is to keep the unwanted out, the wanted in, and the cyber miscreants down. According to an interviewee, a community becomes useless once it reaches a certain number of people (Interview 82).

The main functions of trust within the security community are to ensure its effectiveness, to keep the antagonists out, and to safeguard the communities' ecosystem. The effectiveness of the community rests on the members' skills and appropriate roles in organisations to implement response measures. Keeping the antagonists out hinders them from disrupting or disturbing the community and its services. The community's functionality requires an ecosystem that nurtures technical experts' motivation to contribute, their affiliated organisations' willingness to let them spend paid time on community work, and deters appropriation of community services, e.g., for competitive purposes.

In security communities, trust is reflexive security. Mutual trustworthiness is the internal glue against the antagonists and ensures the functionality and effectiveness of the community's organisational security architecture.

Secrecy towards non-members enables community members to openly share information and other resources with their peers. In a way, secrecy to the outer world is a prerequisite for openness within the community, even though sharing within the community still follows the need-to-know paradigm. [132]

[132] One could propose that openness is higher for those resources within the community, which are less of a risk for sharers and are less of a vulnerability by a potential misuse. This however would require more insight into the communities.

7.3.3 Governance techniques supporting trust and secrecy

A good example of how people become members and gain the trust of their peers is provided by one of the younger members of one of the core groups. During his graduate studies, he used to log on some of the more open IRC channels on security. After a number of high-level discussions he was invited to another IRC channel, where he could hang out and chat with some respected members of the security community. However, he was side-lined once again when he, in adolescent inquisitiveness, asked for a source code of a software used by cyber criminals on the IRC channel. Eventually, he set up his own honeypot, captured malware and dissected it. Now he could offer that source code to the community.

> "Once I showed that clue, that level of understanding of technologies, of the security world and the ability to conduct an operation to get that kind of information, I started becoming an insider, started to become trusted." (Interview 21)

In addition to his skills, he had been in the community for a long time and the community had checked his background — he worked for a respected organisation — in the meantime. Becoming trusted here means that he was invited to private channels "where there's already some level of trust, and we could talk about a lot more operational things" (Interview 21). Later, he developed a technology useful for the security community, which got the attention of another trusted researcher in the community.

> "And he started introducing me to his friends and said, hey, I've been working with this guy for several months now, I've been mentoring him, he's now got the flavour and the socialisation or the introduction, the indoctrination into the way we treat things." (Interview 21)

Communities reduce the risks of their collaboration by applying various governance techniques and best practices. One example is the overall adherence to the need-to-know paradigm. Attempting to reduce the risks of potential data leaks, persons who want to share certain data with others often do not send the whole data package to the entire list. Instead, he or she[133] would only share samples or describe the data to be shared, so that anyone who assumes to be in a position to act upon that data would need to ask for details. Alternatively, the sender would try to identify appropriate addressees and directly send the information to them.

[133] Only some 5% of my interviewees from the technical community were women.

A second example is the non-delegation of trust. As mentioned above, communities try to reduce the risk of admitting unhelpful outsiders, whether malevolent or benevolent, by vetting candidates. The accession decision then lies with a few persons vouching for the candidate. Some communities rely on a unanimous vote for the initiation of the candidate. From the perspective of the community, there is the risk that vouchers take vouching as lightly as, say, rating agencies have looked at CDOs. To reduce the likelihood of vouching being lavished too lightly upon new candidates, vouching members are held responsible for the activities of a member they have vouched for. The responsibility is not transferred from the voucher to the entire community once a candidate has been promoted to a fellow member. The consequence of a member losing the trust of the community is that his vouchers might likewise lose the community's trust.

While vouching (cp. section 6.2.2) appears at first sight to be not more than issuing some moral support, it also is a risky act for the warrantor. By vouching for a new prospect, the warrantor links his credibility within the community to the prospect and his future behaviour within the community. A new member could harm the community and the vouching person. He could appropriate the information and internal services of the community; compromise the sender by publicising private, sensitive data another member has shared; lure other community members into self-inflicting actions, e.g., by sending them faked lists of allegedly harmful machines; send information about ongoing investigations or imminent botnet takedowns to respective malefactors. A loose cannon has the potential to endanger the output of the community's work and nullify months of effort. A community member will therefore presumably only vouch for a new aspirant if a deep trust relation already exists between them.

An important factor for characterising the Internet security community is the role of companies for the community. While all members apparently work in IT or computer science, only a few members' day-jobs are not related to the community work at all. Most members of the community are affiliated to organisations that are somehow affected by Internet insecurities or are involved in providing enhancements to Internet security. Unsurprisingly, the interest of the employing organisation is not overlooked by many community members.

> "And when we're in the security community, we're supposed to give as much indication as possible as to which interest we're acting in, but try to act in the interest of the Internet as often as possible. We all have a boss, we all have to pay the bills, so we all have to do things in the interest of our companies." (Interview 21)

Just as individual companies benefit from the community, the community seeks the inclusion of major ICT players. Despite corporate interests, the glue of the community is interpersonal trust. Even larger players in the field usually have only very few persons serving as community members, thus serving as liaison between their home organisation and the security community (Interview 86).

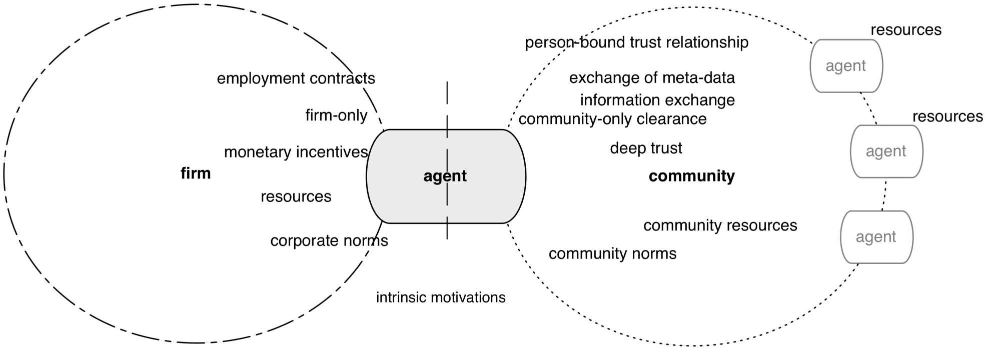

Figure 7.1: The expert as broker between firm and community

Even though access to the community partly relies on the resources controlled by community members in their day jobs, communities are essentially driven by interpersonal rather than inter-organisational rules. The norms within NSP-Sec, an operational Internet security community that is probably backed by the industry like no other, highlight the elevated role of the individual in these communities:

> "People in NSP-Sec have to be able to trust that they can share information with other people in the group that will not in turn be shared with those people's employers. They need to be able to ask for a takedown from one of their competitors that know that that takedown will occur without any further evidence or argument. That's a very high level of trust. The level of trust between people within the group is higher than with any of their employers or co-workers." (Interview 23)

Only when loyalties towards the community are stronger than to the employer or other external entities, members can mutually trust each other that even peculiar information is not misused at the cost of the sharer. This distribution of loyalties ensures that members act in favour of their communities' goals and not simply reap the fruits of their collaboration with the community. This distribution of loyalties is disturbed when communities or single members face strong external demands, e.g., by politicisation or employers bullying their list members to make an exception and be less restrictive with the data, which has been shared with him. The community's internal architecture and governance structure are at risk when such demands arise.

7.4 Rationalising the organisation of Internet security production

The hindrances to openness have been discussed in the previous sections. The security community's need for secrecy has led to restricted membership and sharing policies. To understand why certain elements of peer production can be found in the security production model of the two cases, a deeper look at the inner workings of Internet security production through the lenses of peer production theories is required.

The following subsections build on section 2.1.5, Feasibility of peer production, which discussed arguments regarding the feasibility of peer production as an alternative form of production. Applying the categories and characteristics developed in this feasibility model, the subsequent sections then portray the realities of security production, which have been described in previous chapters and sections of this study. Section 7.4.1 compares the incident response economy with the defining characteristics of the networked information economy. In the subsequent section, this juxtaposition is deepened by an analysis of the particularities of motivations to contribute to security communities. Furthermore, the economic and governance aspects of peer production are compared with those of community-based incident response. This subchapter aims at finding the answers to the question why certain elements of peer production can be found in the cases analysed earlier.

7.4.1 The Internet security production economy

The Internet security economy differs slightly but decisively from the networked information economy. While most of the characteristics described in section 2.1.5 can also be found in the security economy, there are some decisive differences.

Creativity and human intelligence appear to play a similarly pivotal role in the Internet security economy than in the information economy. In the domains depicted in the two case studies, creativity has trumped capital investments. Evangelists of new production methods, for example, in journalism used to mention that one no longer needs a publishing house with huge printing presses and a countrywide car- and train-based delivery network to create news. In a similarly oversimplifying fashion one could state: security production on the Internet does not require insanely costly and proprietary air- and land-borne reconnaissance and defence systems, but first and foremost clever minds who know how to architect, configure and program systems to keep intruders out.

Human creativity is certainly at the core of responses to Internet security incidents. However, given the large number of overall incidents and the repeated application of similar attack patterns, incident response also is simple, plain, uncreative, repetitive work. Employees of organisations who have to defend their networks several times a month against DDoS attacks, are unlikely to feel creative after the tenth such response. That leaves room for more capital intense automation of tedious tasks. Creativity however is required when new things occur, when unchartered territory is entered, either by the attackers or from the response side. Then certainly human creativity trumps everything else.

Second, capital costs are rather low for anyone willing to contribute to the production endeavour. Capital costs required to establish community-based collaboration is nothing to speak of. However, there are some costs involved for hardware and bandwidth for some aspects of security production, primarily sinkholing. This is comparable to large-scale peer production projects that need to host and deliver content to huge numbers of users.[134]

Third, direct capital costs likewise should only accrue to negligible amounts for exchanging information, having discussions in these groups, thinking, coordinating and other purely mental work. Some important activities in the response process however require substantial capital investments. Much of the non-profit services provided by the Shadowserver Foundation run on serious hardware that is capable of digesting all the telemetry coming from honeypots or similar sensors, which themselves often imply substantial hardware costs. The exchange of information and data however, which is the core feature of the communities, is of as low capital-intensity as for any other economic sector on the Internet. This hardly comes as a surprise as the Internet security economy uses the very same communication technologies as the networked information economy in general: the Internet.

[134] Benkler realised the costs potentially involved in integrating and delivering huge amounts of information and data. His liberal definition of peer production allowed some degree of hierarchisation that favours the integrator in exchange for his investments (Benkler 2002, pp. 436-446). For the uncompromising school of peer production around thinkers like Michel Bauwens, peer production needs to be entirely commons-based and doesn't allow any kind of proprietarity and exclusivity (Bauwens 2012a).

Fourth, information in the Internet security economy is semi-rival, whereas information in the networked information economy is non-rival. For peer production to be viable, Benkler asserts that the “pre-existing information/cultural inputs”, which are then modified in the provisioning processes of peer production, are “public goods in the strict economic sense of being nonrival” (Benkler 2002, p. 377). Non-rival does not mean that it could not come with a price tag. In fact, the core products of our very own production domain, academic research, the peer review of which is given as an example of peer production by Benkler (2006, p. 323) creates journal articles and books which usually are proprietary, costly, but, indeed, non-rival.

The domain of information security differs from that by creating semi-rival goods. The most obvious example is the market of zero-day vulnerabilities. They are purely informational goods that can be copied and consumed without being depleted faster by any marginal consumption or use. However, if a rival learns the details of a zero-day vulnerability in an agent’s possession, it no longer increases that agent’s relative power position. The information itself may not be rival, but the advantage that can be drawn from exclusive possession is.

There is a similar rivalry of information regarding defence mechanisms and strategies used by the security community. Again, the information itself is non-rival, but the effects of exclusive knowledge by the security community alone are nullified when the information is transferred to the attackers. Within the community, some information is rivalrous in the sense that the information is appropriable and some first-mover advantage might exist. For example, companies have incentives to be the first to report a new incident, attack, a new botnet, or news about culprits to achieve best PR effects. This rivalry is partly mitigated by the community's rules and governance practices, but to some extent, sharing is thwarted by the possibility of such unilateral appropriation.

The rivalry of vulnerability data is not a characteristic of that informational good per se, but is contingent upon the social relations between agents with or without access to that good. Among friends, distribution of vulnerability information is a minor risk, if at all; among foes, it is potentially damaging. The value is not in the access to that informational good, but in the exclusive access. Such social contingence of some information is ignored by supporters of the “well-known techie-activist rallying cry” (Wagner 2003, p. 999) ‘information wants to be free’ (Clarke 2000). For many this slogan has prevented the emergence of a more sober view of the realpolitik of information.

To sum up, the main difference between ideal-type peer production and security production is the (semi-)rivalry of some information shared by the security community. Therefore, informational input resources are not entirely open and accessible for anyone. The three remaining variables, the importance of human creativity, and low, distributed capital costs come in similar manifestations.

7.4.2 Motivations

For peer production to thrive, it needs a favourable economic and technological environment (cp. sections 2.1.2 and 2.1.3). Within such a nurturing environment, the information opportunity costs of peer production as a mode of organising production will add up to figures below those of traditional firms and hierarchies. The breeding ground consists of supportive motivations and transaction costs — or in greater detail of favourable gains of information processing, allocation efficiency due to the absence of property, sufficient motivations of contributors, modularity and granularity of tasks, low cost of integration, and unfavourable costs of a property system (cp. section 2.1.5).

This section aims to shed light on the conditions existing within the incident response environment. However, none of the aforementioned variables have been operationalised in chapter 3, nor have the interviews been designed to explicitly cover these topics. Therefore, answers necessarily are anecdotal mixed with some theoretical considerations.

Motivations of contributors have apparently not been a problem at all. The modularity of incident response and the granularity of the respective modules have apparently been diverse in both cases, just as they should be.

> "A project that allows highly motivated contributors to carry a heavier load will be able to harness a diversely motivated human capital force more effectively than a project that can receive only standard-sized contributions." (Benkler 2002, p. 436)

While some persons spend days and nights for weeks to solve certain issues and were absolutely fine with doing so in a voluntary manner, others just posted a few hints or minor details on, e.g., malware involved, attacking systems, or potential perpetrators. The deep intrinsic motivation of those involved in the response endeavours has been described in previous chapters. These persons were driven by convictions, the prospect of indirect appropriation, and other forms of intrinsic motivation, which made them spend considerable time on the response endeavours.

There is one caveat though that has not affected the response endeavours analysed, but it might affect incident response activities in the future. The ability to modularise the overall product into smaller components, which can be produced "independently and asynchronously" (2002, p. 434f), allows for finer granularity of the modules. Fine granularity implies that the time and workload necessary to produce a specific module is lower than with coarse granularity (2002, p. 435). Fine granularity therefore offers opportunities for those only willing to contribute small amounts of personal efforts. Once a collaborative project manages to incorporate very fine modules created by a large numbers of contributors with only minor motivation, it can at least theoretically produce larger products. This is the catch for the security community. Given its reasonable preference for walled secrecy and its existing main organisational techniques to reduce the risks of defection — deep trust, access control, vetting — the security community cannot easily open up its resources and communication channels to the anonymous masses that contribute to other open source or peer production projects. As of now, this organisation path is closed for the security community. Should it need to increase its capacities, it must turn to other organisational approaches. In two cases analysed in this research project however, the lack of finely granulated modules has not been a deterrent to highly motivated contributors.

In theory it might be possible for the community define finer granulated modules, open their production up to the public at large, and thus attract large numbers of occasional contributors. To accomplish that, the community would, however, lose one of its greatest advantages: the relatively low set-up costs of the production process.

7.4.3 Transaction costs

The transaction costs of an organisational approach comprise a number of variables. This section focuses on set-up costs, information processing, and integration.

Set-up costs. An explanatory approach for the existence of some elements of social production in incident response lies in the set-up costs and marginal transaction costs for non-social production frameworks. From an organisational and institutional perspective, effective incident response is a problem of gathering data, analysing information, formulating a response strategy, and coordinating activities of numerous actors world-wide, reconfiguring ICT systems in diverse regions, jurisdictions, and under distinct ownership.

Market-based production requires a high degree of crispness of the products and services to be delivered.

> "A market transaction, in order to be efficient, must be clearly demarcated as to what it includes so that it can be priced efficiently. That price must then be paid in equally crisply delineated currency." (2004b, p. 315)

It takes time and substantial resources to gather the information required to exchange security related informational goods via prices on a market system. In an ideal-type market for such Internet security incidents, all the aforementioned tasks and information exchanges would come with a price tag.[135] In theory and given the technical nature of the attacks, any attacked person or organisation might rely on the cooperation of any organisation or person with systems connected to the Internet. On a market-based approach, information exchange can be based on contracts or ad-hoc transactions. This impossible number of potentially necessary relations to exchange information could be reduced by installing a smaller number of feasible proxies such as ISPs, CERTs, vendors of security software or operating systems. But that would still leave us with a probable five-digit number of potentially required agents. Again, in case one would want to base exchanges of information and knowledge on a network of bilateral contracts between agents, the number of possible relations among these agents is still infinite.

The alternative to contracts is the exchange of, for example, attack data on a respective market for such data. It requires little thought to see the problems with such a market. The most obvious hindrance is that the supply-side would structurally be at an advantage as the demand-side urgently needs any information regarding an ongoing attack and would be forced to buy-up large parts of, if not the entire supply, regarding a specific attack. In addition, all the trust issues with sharing described earlier in section 7.3 would persist. Whether the problems that are inherent in such a market can be solved by appropriate regulations or not, requires deeper analysis. It requires substantial upfront set-up costs to establish such a mar-

[135] Admittedly, organisations have increasingly developed trust-based forms of exchanging valuable information and service in so-called value networks (Christensen & Rosenbloom, 1995). The characteristics of the latter however is that these networks are generally initiated as a result of managerial decision, and are often embedded in some kind of contractual agreement among participating organisations. No such thing exists for the security community.

ket in the first place. The market for zero-day exploits (Egelman, Herley, & van Oorschot, 2013), which likely is substantially smaller both on the supply and the demand sides, serve as an example of the complexities, advantages, and drawbacks of markets for security information (Stockton & Golabek-Goldman, 2013). Regarding attack information supply, the Conficker response community set up its own honeypot and sinkhole systems; the Estonians received relevant information within their Estonian community and later via global security communities.

Table 7.3: Hindrances to social production (left column) and their circumvention in peer production (Benkler 2002) and community-based Internet security production

	Peer production	Internet security production
Lack of tasks for low-motivated contributors	Modularity and granularity of tasks to skim low-level contribution	n.a.
Defections	Redundancy by high # of contributors and modularity; averaging out outliers	Trust-based access control regime; repeated vetting; chained 1-strike rule
Free-riding	High # of contributors Contributors' motivation not affected by # of users	No–lurking rule
Unilateral appropriation	Copyleft licences	No-appropriation rule and expulsion sanction
Integrity assurance	Not required (Benkler); graduated introduction (Dafermos 2012)	Access control; chained 1-strike rule; graduated introduction
High integration costs	Lowering costs by: automation Dealing with costs: integration by markets/hierarchies; cooperative appropriation	Dealing with costs: donations by industry; self-financed (expecting high-yield indirect appropriation)
High lossiness	Open communication	Network of communities linked by overlapping liaisons

Security production, however, comprises much more than the exchange of attack data. To understand the problems involved with a market approach, it is worth considering the implementation side of any solution. Assuming that the response community had found an appropriate way to respond to an attack, one or more agents would have to buy the services of other agents who implement the required reconfigurations on relevant systems. The issue is: who could be that buyer? Using the example of Conficker, it is hard to imagine who could have bought: the service to update all Windows computers world-wide and then installed the required bugfixes onto them; the response team to coordinate and manage the entire response endeavour; the trust necessary to facilitate close collaboration in the first place. The point here is again not to assert that such a market approach, especially when combined with some legislative regulation, is not feasible, but to argue that the establishment of such a market involves high transaction costs and may require tremendous upfront investments.

The problem of substantial set-up costs, albeit presumably on a lower level, also holds true for hierarchical provisioning of Internet security. The obvious problem with an approach is that it is hard to conceive of which agent should gain sufficient authority to address such Internet security incidents. The degree of authority necessary for an ideal-type hierarchical approach is substantial. It would include the ability to monitor any system connected to the Internet, the ability to reconfigure these systems, the ability to access the source-codes of any commercial software for analytical purposes, and much more. Clearly, such a system is neither normatively desirable nor politically feasible. Even modest steps towards a more hierarchical approach could take years to be formulated and pass through the legislative due process.

While the costs to set-up a market or a hierarchy for incident response would be substantial, a global security community does not emerge out of thin air. The communities' prevalent form of trust, deep trust, takes months if not years to establish between individuals. It usually requires repeated collaboration over a longer period of time to become a member of the core communities. Setting-up and perpetuating a community based on deep trust, vetting, and access-control requires substantial attention by the contributors. The set-up costs would have been even higher, if the community had not opted for a walled-secrecy approach by default.

Information processing. A second major variable in the transaction costs formula are the costs and gains of information processing.

Benkler's model of peer production viability has a variable that circumscribes much of what fuels the necessity for trust and secrecy given the presence of the antagonists: the costs and gains of information processing by a production system (Benkler 2002, pp. 407-415). To take appropriate action, a mitigation response organisation needs adequate information and knowledge resources.

Any organisational approach brings its own degree of "lossiness", i.e., the difference between "perfect information" and actual information. The price for this difference results in "information opportunity costs" (2002, p. 408). Lossiness occurs either because information is distorted on its way from the source to its target or arrives too late. Information opportunity costs also arise when a person with valuable knowledge and information does not participate in a collaborative effort or valuable information is considered in the production process, or when valuable information is not shared because no other adequate coordination mechanism is in place for that particular information.

While the separation of the communities might hinder the influx of valuable information, a theoretical comparison with other modes of organisation is in favour of the communities' walled-openness approach. The plausibility for the trust-based sharing approach becomes obvious by comparing it with the coordination mechanisms price and hierarchy. Both require a great deal of knowledge about what is going on, what is shared, who needs information, etc.

Moving all the "clusters of resources and agents" (2002, pp. 372, reporting Coase's definition of the firm), which interact in the community based trust, intrinsic motivation, managerial backing, to another organisational form seems like an impossible task. The task of Internet security in general and incident response is too huge, too innovative, too fast-paced to specify products and prices in a market or work-packages and responsibilities in a hierarchy for all aspects of Internet security provisioning. What's more given the current technological state of attacks, Internet systems and defence approaches, the range, distribution of required actors is simply too large to incorporate the network into a hierarchy, be it a Global InternetSec Inc. or a United Nations Global Peoples Internet Security Organisation.

In the security world, another term needs to be added to the equation of costs related to information. With the presence of an antagonist, the risk exists that someone uses information produced by the community in a way that undermines the community's efforts or imposes costs on a sharer of sensitive information. But it's not just in the arena of cybercrime, where actors are confronted with antagonists that are willing to use any information, product, or service provided by their opponent, irrespective of the interests of the originators. As the saying goes: All is fair in love and war. And so, alas, in realpolitik. Therefore, as soon as the field, in which a community is playing or their products are located, is politicised and fiercely contested, the information gains/costs variable turns against the viability of peer production. The formula for information costs needs to incorporate the risks information distribution.[136]

Consequently, the prerequisite for openness in a field containing antagonists is hard work. It requires the definition of information that could be shared with a

[136] This is not to be mixed with a tragedy-of-the-commons game, in which the coherence of a community falls apart as this game supports the defector. It is a community that plays a game in which it either wins or loses, and information dominance is the variable that decides about the outcome of the game.

public audience. Hence, the set-up costs for a peer-production regime require substantial ex-ante investments. It requires a painstakingly accurate analysis of what is shared, and the risks that could be attached to each information item, to finally recommend what should be shared only one-to-one, what should be shared within trust groups, what should be shared within a larger community and what should be shared in public forums. Even then, the risk remains that the antagonists could exploit released information items in an unforeseen way. The openness of the defence side ironically allows the antagonists to use peer production himself and enjoy the disruptive innovations that can emerge when persons from various backgrounds find unforeseeable combinations and ways to circumvent or neutralise the work of the defence side.

Integration. A third major variable to explain the low transaction costs of social production are advantages in the integration of diverse contributions. Motivation, crispness, and set-up costs mostly linked to the problem of how to ensure the required influx of contributors to a voluntary production regime. As it is based on intrinsic motivations, it needs to avoid any factors that sap contributors' said drivers or otherwise weaken the crucial provisioning and integration phase of the production process. Benkler has identified a number of potential pitfalls and has discussed ways in which to circumvent theses hindrances in peer production projects. A key concern that Benkler has subsumed under the integration task are the costs of defections. In the ideal-type peer production model, the occasional under-performer and the malevolent vandals are made up by high numbers of contributors. For the security community, the risk of defection is substantially higher as the likelihood and the costs of damage are higher. Therefore, the community has applied a variety of governance techniques, among them the no-lurking rule, access-control, the chained 1-strike rule and, akin to the FreeBSD community (Dafermos 2012), a graduated introduction of new contributors.

7.5 Conclusion: A variant of social production

This chapter has aimed to shed light on the question of why some elements of peer production have been applied in the cases studied in this research, while others have not.

The section "Drivers of secrecy" revealed that the antagonists or "bad guys" have been the main driver for the community to opt for a walled organisational approach. Secrecy towards non-members is a guiding principle while fostering open communication within the confined borders of the community. Surprisingly, states' activities have at best played a marginal role in the closure of the cases' response teams and the security community in general.

Given the frequent mention of trust as a defining characteristic, the community's stance towards trust and the respective governance approaches linked to that concept must be explored. The response networks and the community in general are overall based on deep-trust relationships. The importance of trust within the community problematizes Benkler's purely economic approach. In his *Coase's Penguin* essay, Benkler builds a model of peer production based on institutional economics thinking, leaving *trust* entirely aside. Trust has no place in Benkler's theoretical model of peer production (Benkler 2002) or social production (2006). However, the model has included a number of variables that are either closely related to trust, are prerequisites of trust or are the cause for the need of trust. The application of deep trust as a prerequisite for community membership can therefore be interpreted as a hedge against the significantly increased risk of defections in the shadow of the omnipresent, yet invisible antagonists.

The last section of this chapter argued that the economic and technical conditions of security production in the cases analysed in this study justify the organisational approaches of the responding communities. The set-up costs of a non-market and non-hierarchical approach are still below the investments required for an approach not based on a trust-based network. The motivations of contributors have been sufficiently high to contribute to non-fine granulated modules. The fast-paced changes in the Internet insecurity markets requires similarly flexible responses by the defensive side, and networks are known for their flexibility to adapt to external challenges. Along with this, there have been plenty of external demands that help to explain the organisational departure from an ideal-type production model, and the adoption of an altered variant of networked production. By adopting them, the Internet security community has ensured its ability to collaborate with distributed actors based on social relations as a guiding organisational principle.

On the phenomenological front, the decisive characteristic of peer production is openness, whereby openness implies easy access to input resources; openness of the production platform and ease to join that process; and finally the unhindered accessibility and reusability of the product created in the production process. In addition, individuals participating in the production process and use of resources are not steered by monetary incentives and hierarchical authority as organising principles. Instead, peer production rests on intrinsic motivations that let contributors spend time and personal resources on a collaborative production process.

Neither the data nor the research approach of this study allow for exact statements about the causes of the actual shape of the production model. One could however hypothesize that observed instances of the social production model resemble an institutional design that attempts to incorporate some major advantages of the ideal-type peer production/open source model, while at the same time factoring in the need for secrecy. In terms of transaction costs, the observed hybrid, not-quite-peer-production flavour of social production reduces the risks of intrusion by malevolent bad guys, who seek to nullify the communities' defence efforts. On the other hand, it keeps transaction costs of secrecy relatively low by using community-based vetting procedures and remaining somewhat permeable to new contributors. The open source access-for-all policy and the swift trust model, which is based on assumed common values and goals (Adler & Heckscher, 2006, p. 24; Osterloh & Rota, 2004, pp. 13-16), is replaced by a system of limiting access to its infrastructures, vetting potential contributors and sharing on a need-to-know basis only. While secrecy thwarts one source of peer production effectiveness — the unplanned, unrestricted use of resources by high numbers of agents with diverse talents and skills — security communities can still leverage relatively low-cost permeability to new contributors to take advantage of external "information gains" (Benkler 2002, pp. 407-423).

8 Conclusion

"But in republics there is more vitality, greater hatred, and more desire for vengeance, which will never permit them to allow the memory of their former liberty to rest; so that the safest way is to destroy them or to reside there"
Nicolo Machiavelli[137]

This study has investigated the responses to large-scale Internet security incidents. By exploring their organisational approaches, it has aimed to analyse whether and to what extent the ideal-type form of networked governance, peer production, is applied in security governance.

Peer production is regarded as an innovative form of production that has been facilitated by the rise of information technology, the Internet, and their diminishing effects on transaction costs for the production of intangible goods in a distributed, collaborative, self-organised way. Existing research and theories on open source and peer production have focused on describing the enabling factors of this mode of decentralised, distributed, and non-hierarchical form of collaboration. Yet the limits of this form of production have not yet been tested, nor have the organisational hybrids been explored that might emerge when some of the theoretical prerequisites for peer production are not met. This study therefore aimed at contributing to the literature on open source/peer production by analysing its limitations and feasibility in the domain of security governance. The second research gap this study aimed to address has been our sparse knowledge about the actual practices of Internet security production.

Two underlying fundamental societal issues are at stake here with these research goals. One raises general questions about the possibilities of innovation in security governance; the other more specifically affects the institutional design of Internet

[137] Machiavelli 1505/1910, "Chapter V — Concerning The Way To Govern Cities Or Principalities Which Lived Under Their Own Laws Before They Were Annexed".

security governance. These issues set the context for this study. The first is the feasibility of peer production in domains other than open source software. Openness and transparency have been lauded as the cure to countless societal ills. At the same time, traditional security institutions are known for their restrictions on accessibility, democratic control and supervision. Their traditional secrecy, their reduced accountability combined with the exertion of the monopoly of force, lets openness appear as a panacea to their problematic democratic legitimisation. The second is the Internet's institutional polity, the design and nature of the institutions developed and chosen to govern and maintain the Internet. At the starting point of this research project, Internet security policy discourses were in their infancy, and the question of which governance approach for Internet security will serve our societies best is still far from being answered. Certainly, the complex questions raised above go beyond this study and will continue to persist. This is the bigger picture to which this research project has aimed to make a detailed contribution to. In comparison to this wider context, the actual research questions and the actual research design of this study have covered a more modest area, encompassing only a tiny aspect of the field.

The field of Internet security governance was already in flux when this project began in late 2008, but in the years that followed it has turned into a media frenzy, a political hurricane, and the contested field of great power politics. Incidents, conferences, media frenzy, policy drafts, legislative programmes, and international conflicts have followed in quick succession. The list of incidents discussed in relevant media outlets and among security pundits is stunning: Anonymous, Al-Qassam, various botnets (e.g., BredoLab, Coreflood, DNS Changer, Grum, Kelihos, Koobface, Mariposa, Nitol, Rustock, Storm, Waledac, ZeuS), the Brazilian Sat-Hack, Cablegate, Comodo, DigiNotar, Duqu, the Google-vs.-China strife, Tunisia, Egypt, and the Arab Spring, the Flame, Gauss malware, cyberattacks on Georgia and Lithuania, HBGary hack, Lulzsec, McColo, MiniDuke, Operation Shady RAT, Operation Ghost Click, Pakistan's YouTube Re-routing, Red October, South Korea cyberattacks, TurkTrust, VOHO, attacks on the Washington Post and other US organisations, the watering-hole attack on the website of the Council of Foreign Relations, phenomena like Wikisposure, Wiper; and of course Wikileaks, Stuxnet, and the Snowden leaks on the massive surveillance of Internet traffic and services by Western intelligence services. The policy debates about national cyber-strategies in various countries, ACTA, data retention, the Internet kill switch, net neutrality. Countless security reports from many security solutions providers; the rise and maturing of national institutions for public-private cooperation; the formulation of cyberwarfare doctrines; web-filtering of child-abuse imagery, bomb-construction manuals, porn; CISPA in the US; the emergence of Internet content as national security threat in authoritarian countries; the rise of a market for zero-day vulnerabilities; the maturation and development of high tech crime teams in LEAs; the creation of national cyber-boards in a number of countries. This is all further complicated by the involvement of a wide number of national governments, international organisations like Council of Europe, Interpol, ITU, NATO, EU; initiatives like the London Action Plan, the Messaging Anti-Abuse Working Group, ICANN; the security industry. There have been few, if any policy fields that have been more in flux than Internet security in the last couple of years.

This research has only covered a relatively small section of the overall Internet security empirics. Nevertheless, the analysis of these two cases observed through the lenses of peer production and open source theories provide some interesting insights into the intricacies of real-life Internet security production. On the methodological front, the study provides a coherent operationalisation of peer production. The narration of the Estonian attacks and the Conficker botnet, and their respective political and technical circumstances, contribute to the emerging field of the historiography of Internet security attacks and incidents. The depiction of the response activities provides one of the first accounts of how Internet security production and global large-scale incident response works, especially on the operational level. The study identifies ad-hoc teams and Internet security communities as decisive actors in response endeavours. It provides details to their coordination mechanisms, and the first scientific analysis of this community and its characteristics.

Noteworthy from the perspective of a student of governance systems is the observation that distributed, bottom-up collaboration systems can play a substantial role in the mitigation of distinct incidents. The narratives of responses to attacks describe how distributed, collaborative security production works. The response process cannot be categorized as peer production in its narrowest sense, however. This did not come as a surprise, precisely because some degree of secrecy among the response teams was expected. More surprising however was how much these response activities, the inner organisation and the exchange of information, followed existing policies and relied on a standing Internet security community. The response was not just, as assumed in the starting phase of this research project, the result of a "loose network of technical experts", who joined forces when their skills and influence were needed. The response teams were formed in an ad-hoc manner, yet they relied on existing networks of trust between the actors involved. Their collaboration has furthermore been facilitated by the interplay of several established Internet security communities with a common set of cultures and rules, and different mostly technological focuses. Interestingly, there were not substantial differences in the overall organisational approach of the response in both cases. The response to the Estonian attacks did not come from traditional security organisations, despite all the cyberwar rhetoric by Estonian politicians, but from a collaboration of the Estonian Internet security community, in which intelligence services and police only played a minor role, with different global security communities. A lack of competence and coercive technical authority, organisational precautions and a lack of preparation for such incidents are among the reasons why traditional organizations did not meet the challenge. The Estonian case also highlights the importance of global collaboration and information exchange to allow the network of security communities to act as the global Internet security task force. The Conficker case on the other hand highlighted the strengths and weaknesses of the community approach. The response efforts have been a show of capabilities by the security community in general and the CWG and its affiliated actors in particular. It showed the technical finesse, the responsiveness, and the organisational skills of the community. At the same time, however, the hardly sustainable amount of work put into the project by contributors, the financial burdens of some contributors, and the spectre of simultaneous attacks in the Conficker class give an idea of the limitations or at least weaknesses of the community approach.

The following sections summarize some of the findings and contributions of this study to the existing body of literature. The first section picks up the question of whether and which elements of peer production can be found in existing Internet security production endeavours and the effect it has on security production. The next two sections add remarks concerning the limitation of this research and suggest future paths for research and further development of the ideas. The last two sections then remove the rigid methodological corset of this study and consider the state of Internet security governance and ways to enhance the distributed approach. Section 8.4 aims to contextualise the findings of this study in the wider arena of Internet security, its production and governance. The subsequent section finally puts all scepticism regarding the future of the egalitarian, networked approach aside and discusses ways to improve the latter.

8.1 A different way of producing security

The previous section listed a number of incidents and topics in Internet security that have received substantial public awareness. Despite all the societal focus on Internet security, there has been little research on actual security production. The public debates overlook the question of whether new forms of collaboration facilitated by the rise of the Internet can be used for the very protection of the latter. Could collaborative, open, accessible, egalitarian production projects possibly replace or at least supplement hierarchical authority and the market's pecuniary incentives to achieve societally beneficial outcomes? How far can this new thing called open source production or peer production be pushed? It has been unclear whether the need for secrecy, a common characteristic of security production, would inevitably break openness, which in turn is a decisive characteristic of peer production. One could argue that security institutions, their centrality, and the mode of governance of coercive, institutionalised force are defining characteristics of any society.

Forms of peer production existed in the wider world of Internet security — that was apparent at the outset of this research. Phishtank.org and the Whois database[138] were mentioned as an example of the general feasibility of peer production to produce the informational resource necessary for incident response. But aside from these examples, does peer production really matter in serious Internet security issues? The best way to find an answer to this puzzle is to look at some serious large-scale Internet security incidents.

The first research question that has guided this study is whether the handling of a particular Internet security incident can be classified as peer production, or at least as predominantly peer production. The second set of research questions were conditional, depending on the outcomes of the first. Given that the incident responses could not be classified as exclusively or even predominantly peer production, the question then became whether there were still substantial elements of peer production used in the response to these incidents; which elements were used; for what purposes; and what the primary causes for the non-application of a pure mode of peer production were. These research questions already indicate that the wider issue of "Internet security production" was here narrowed to the realm of responses to large-scale Internet security incidents. Two cases were selected and studies conducted by desk research and qualitative interviews: the attacks on Estonian Internet-connected systems in 2007 and the Conficker botnet in 2009.

To facilitate the answering of these questions, chapter 3 developed an operationalised model of peer production and a model of responses to Internet security incidents. Chapter 4 provided a narrative of two high-level attacks, which put the Internet's functionality at risk. The following chapter then described the response activities in the two cases studies. These two chapters set the scene for Chapter 6 on the Social Dimension of Internet Security, which answered the question as to what extent the incident response was peer produced, and Chapter 7 on the Limits

[138] There is a discussion to replace the existing database with a new version, to which law enforcement agencies would get access to, but not mere Internet users (Kuerbis 2012).

of Openness, which explored the reasons why some elements of peer production could be found in the organisational approach of the response in the cases analysed.

But first, peer production and security production needed to be operationalised to allow for the breaking down of a given production process into its component parts to analyse its ‘peer-producedness’. Peer production in its most rigid commons-based flavour is defined by decentralised collaboration, unrestricted information sharing, non-proprietarity and reusability of produced goods, collaboration and sharing not based on market mechanisms, and finally collaboration among actors not based on hierarchical commands—or in short: distributiveness, openness, and socialness. These defining elements of peer production were themselves detailed by a number of identifying characteristics, such as a particular set of motivations for contributors. Despite these straight criteria, it required at times somewhat arbitrary decisions, e.g., on whether particular tasks within the response process were based on hierarchical commands or socialness.

The answer to the research question of whether the handling of a particular Internet security incident can be classified as predominantly peer production was given in chapters 5 and 6. The former provided a historiographic narrative of the response to the two incidents. While both response endeavours were largely driven by a mix of voluntary collaboration, one can argue about the corporate backing. But even if one would deem some contributors as ‘backed by their company,’ it does not fundamentally change the characteristics of the collaboration. It is a pattern familiar from open software communities such as Linux or LibreOffice (Vignoli 2013) that the majority of commits and contributed lines of code are created by employees of companies, for which the open source software is a necessary component for one or more of their product offerings. Open source communities have established governance precautions against co-option of the production community by large donors, e.g., by establishing thresholds for maximum voting weights for internal bodies. While information sharing and co-collaboration are visible elements of the response communities, some important elements of peer production are missing, most notably openness of or unrestricted access to the production platform and non-proprietarity of produced goods.

The accessibility of the produced outcomes are a close call and depend on the level at which one looks at these things. At the most high-level view, the product of the response endeavours is a re-established and assured level of Internet functionality; i.e., increased security. Such security has the characteristics of a public good that can be enjoyed by anyone and that cannot be over-consumed. At lower and more tangible levels, however, everything that is produced by the response communities usually stays in these communities. Results like code or new settings rarely end up at open repositories like github. Each community has its ‘secret sauce’ that hinders third parties from easily forking the community as the source code or raw data to intermediate products usually are somewhat proprietary.

Another fundamental difference to ideal-type peer production systems is that the platform, on which post-incident Internet security is produced, is not openly accessible and does not own all the means of production necessary as they reside with those who own and operate Internet-connected systems. Access to the production platform is restricted by the communities’ trust-based vetting and access-restriction policies.

To understand Internet security production in its entirety, it is necessary not only to look at the ad-hoc response teams and their individual members, but also at third parties collaborating with responding communities and task forces. In a way, they act like the head of the overall response mechanism, and function as a coordinating layer. The eyes and ears that detect abnormal traffic patterns are at least partly elsewhere, controlled by different actors, for example the members’ employers. The technical and operational authority to actually alter existing systems in ways that increase the Internet’s security reside with these external organisations.

The second research question asked why the incident response endeavours could or could not be categorized as predominantly peer produced, but also why some elements of peer production had been applied. In both cases, the overall response was distributed, and necessarily so, as the narrative of the incidents has shown. The communities of technical operational experts in different geographies and technical fields have been indispensable for the technical mitigation and remediation of the incidents. It is the community members who exert operational control over those systems and machines that are eventually attacked, used for the attacks, or are necessary for the mitigation effort. These community members hold the information necessary to detect, understand, and counter the attacks. In the Estonian 2007 incident, the Estonian and global technical security communities have been mostly successful in their response efforts. Decentrality and technological restraint however has been questioned by some members of the community. At a Dutch security conference in 2010, malware analysts involved in the Conficker response raised the question of whether a different, more centralised and coercive approach was required. At the core of the discussions, was an issue which has since spilled over to policy circles[139], namely whether distributed technical problems such as huge num-

[139] Dutch Minister of Security and Justice Ivo Opstelten proposed a law that would give Dutch law enforcement the right to force-update machines infected with malware or other sorts of malicious software (Koot 2012; Ministerie van Veiligheid en Justitie 2012).

bers of infected computers should be cured by a centralised technical authority with coercive power over distributed technical systems that fail to comply to certain security requirements. The student of realpolitik will likely see the world of Internet security politics and governance develop to a contested place: the realisation of the power potentials has created the power vacuum that actors will want to seize to bar other actors from doing so.

The decisive limiting factor, however, has been secrecy, nurtured for various reasons as the preceding chapter has shown. The "bad guys" are perceived as a significant risk to the success of the community's undertakings and to the privacy and even the personal well-being of community members. The need for secrecy reduces the accessibility of the production platform and to intermediary products, while the overall result, the Internet's re-established functionality, can be enjoyed indiscriminately. While the response endeavours by the security communities partly resemble open source projects, there are substantial differences to said accessibility and openness. Multiple access layers have been set up to shield the community's core assets — information, trust-based relations, and the community's integrity — from malevolent intruders.

Despite these curtailments of ideal-type open source communities, Internet security production in the cases described in this study differ substantially from traditional security production. Most striking is the relative absence of hierarchical structures in the networked security production models. Certainly, there are actors with greater influence or power[140] than others: owners or administrators of mailing lists, those with high centrality and connectedness within a single community, those serving as liaising nodes between different communities, large companies with significant security programmes and budgets for external activities. Nevertheless, none of these actors amasses the degree, let alone the institutionalisation of power, that traditional national authorities have in policing, intelligence, and national security.

8.2 Limitations of this study

The limitations of this study are as apparent as the draughts in an aging barn. This has been a study of two cool ideational kids, peer production and Internet security, running into each other for the first time and seeing how they get along. Limita-

[140] On the synonymous usage of both concepts: Nye 2011b, p. 11.

tion can come by design, by execution, by the nature of the findings and altogether influence the validity and generalizability of the findings (Klievink 2011, p. 227). This section lists and discusses some of the limitations of this study.

It is apparent from the choice of case studies that this study has explored only a subset of what can be called Internet security production. In chapter 2 , Theoretical Foundations , Internet security has been conceptualised as the absence or reduction of Internet-based or amplified threats to acquired values. Re-establishing the Internet's technical functionality, which has been the core achievement of the response to the two cases analysed, is only one way to reduce such threats. It is a response strategy that is relatively close to the capabilities, interests, and every-day tasks of many of the community's members. However, the range of those 'values' has been broadened by various actors from outside the technical community, asking for different response strategies, different security objects, and different response institutions. Responses to incidents in which other sorts of Internet 'securities' are affected will likely yield different results with regard to the question of the peer-producedness of the response. One could hypothesize that more hierarchical, more secretive, more coercive, less transnational responses can be expected especially when Internet security meets national security interests — despite the findings of the Estonian case.

While the Estonian attacks have apparently created a "national security situation", the response itself was driven by the civilian CERT, response teams in major banks and insurance companies, and supported by international Internet security communities. The existence of elements of peer production in this particular national Internet security incident does not allow for conclusions regarding the re-occurrence of these elements in future cases. It does however show that this particular national security situation could be handled by the technical community without the support of military cyber-units. The question as to what extent the distributed, networked response model could be applied to national Internet security incidents in general cannot be answered by this study. Equally, it has not been answered by proponents of more hierarchical approaches to distinctive Internet security organisations. On the contrary, even the DHS funded report described the Conficker Working Group, supplemented by a few institutional changes, as a model for future incident response endeavours (Rendon Group 2011).

The selection of the cases implies another limitation on the generalizability of the findings on the existence of elements of peer production. Both cases took place before a major shift in Internet security policies in many countries and, first among them, the U.S. government (Mueller 2010, p. 179). Since then, the landscape of Internet security and the governance thereof has changed significantly (cp. section

8.4). Thus, the response organisation has an element of historic contingency, coined by transitional circumstances when governments' full awareness of cybersecurity was yet to unfold and Internet security policy was at a relatively early stage. On the other hand, the Internet security community still appears to be up and thriving as interviews in 2012 confirmed. Desk research on responses to post-Conficker botnets indicates that the role of the Internet security community is still vital to botnet response, even though the hierarchical nature of relations within the community appears to have increased in the last couple of years (Schmidt 2014).

Another limitation to the validity of the results is due to the focus on individual community members, and the treatment of their hosting organisations more or less as black boxes. Statements by employed interviewees as to whether they largely acted independently from their employers might have a bias towards community member's independence. A number of questions pertaining to the role of the employers tried to mitigate this effect, and in some two cases claims about the alleged independence and voluntariness of interviewees do not match their presumable job descriptions. There are reasons why an employee might have answered that he acted voluntarily while the organisation's hierarchy, i.e., higher-level management, might have seen the participation as part of the employee's job. The blurring of private and professional affairs in high-skilled professions has been a trend in management and employee-employer relations. Given the difficulties of identifying and measuring reliable performance indicators in innovative and dynamic areas, some employers have abandoned detailed performance requirements and now leave decisions regarding the appropriate means to get the job done and defining the priorities of their positions to the employees themselves. All interviewees have been on the edge of an emerging field of global incident response. Visionary managers could have anticipated the importance of being part of these endeavours and hence have given participants some leeway and a great deal of freedom. On the other hand, the question of employee independence in community contribution primarily affects the contributors' motivations, potentially shifting them from intrinsic to externally-demanded. It does not however fundamentally alter the community's internal characteristics as a networked, distributed and relatively non-hierarchical, non-coercive project.

Related is another caveat, the assertion of the non-marketness of the response endeavours for those contributors who indirectly profited by their participation in the response endeavour. The term indirect appropriation describes monetary advantages that voluntary contributors draw from their unpaid efforts, e.g., by increasing their professional reputation or prominence. Benkler and other open source authors still treat contributions with indirect appropriation as intrinsically motivated. The underlying idea apparently is that personal economic benefits are seen as an accidental outcome or byproduct of otherwise intrinsically motivated actions. Such an interpretation underestimates individuals' ability for up-front long-term investments. Some members of the response teams likely had entrepreneurial ambitions; others might have foreseen an involvement in the Conficker response as a future career booster. The difference between indirect appropriation as a byproduct and as a planned goal is that persons with the latter in mind would not become engaged if there is little chance to reap these advantages in the future. The contributor's calculation might be entirely pecuniary, combined with personal leanings and talents, but in the Benklerian model, it is judged as non-market-based contribution. The actual monetary transaction would indeed only happen in the future. However, one could see the entire package, from unpaid contribution to increased market value to increased earnings as an entirely economic activity. The category of "indirect appropriation" feels like a welcomed black-box that allows the researcher to avoid some epistemological obstacles. In a response to Benkler's depiction in *Coase's Penguin* (2002), Weber criticised the idea of indirect appropriation in open source communities. He argues that companies intentionally develop business models around open source; often, community service or actual OSS coding are an auxiliary service by these companies that are nevertheless necessary for their commercial services around such open source software (Weber 2004, p. 194).

The previous chapter concluded that the 'bad guys' were driving the need for secrecy for the community to protect its members and safeguard the required exclusivity of their knowledge. From a critical perspective, however, accepting the interviewees' explanations of secrecy as a legitimisation for community secrecy potentially perpetuates an ideology of secrecy in security matters that is prevalent in traditional security institutions. Answers indicate that interviewees with a background in law enforcement or with allegedly professional relations to traditional state security organisations stress secrecy more than those with a more academic approach to incident response. This might be caused by different experiences or by a different mindset towards secrecy. Furthermore, the stated need for secrecy does not allow the statement that the current degree of secrecy is inevitable, required, and optimal. It would require thorough risk analyses of information objects involved in the response endeavour to identify the maximum possible degree of openness for the response community.

A letdown for practitioners of incident response and designers of the organisation thereof presumably is that this study does not provide a conclusive, definitive analysis of the effectiveness and performance of the network-of-egalitarian-communities approach. Partial answers to these important questions are distributed all over this study, but the overall research design, the research questions and the methodology do not allow for such statements. But neither have designers of national approaches of Internet security provided scientifically sound analyses proving the overall societal benefits and the superiority of the model of Internet-security-by-national-security-organisations.

8.3 Suggestions for future research

The previous section has given a description of the limitations for the validity and generalizability of this study and its findings. Future research could aim at eliminating these limitations, but also continue to explore the approach of the network of egalitarian communities. The last section in this chapter proposes some practical and more or less feasible means to strengthen the community approach. Academic research could contribute to such a design challenge in several ways. An apparent shortcoming in the cases described was the mediocre interplay between technical communities and national authorities and policy makers. While this study has listed some possible explanations, it is still not scientifically proven knowledge whether cultural misunderstandings, turf battles by bureaucracies and policy circles for their preponderant role in all matters of Internet security, the organisational and institutional superiority of the traditional security organisations, or just plain political interests have lead to the non-optimal relationship between communities and states.

Research on technical Internet security communities in general is only in its infancy. The literature on them is thin compared to other technical communities such as in the software or content-production domain. Many aspects of the communities deserve more scrutiny. This study has been aimed at the response endeavours to large-scale security incidents. The attention given to Internet security communities was a by-product of the investigation of incident response organization. Incident response was examined through the lens of a peer-production model, with its analytical categories of hierarchy, marketness, accessibility, reusability, and secrecy. While these categories have offered a rich picture of the communities, a 'traditional' approach to collaborative, consensus-oriented communities — e.g., based on Ostrom's theoretical models (van Wendel de Joode 2005; Dafermos 2012) — would give complementary insights, e.g., on the sustainability of Internet security communities.

After a plethora of empirical studies on open source software communities, content production communities, and now increasingly of technical Internet services communities, a more general theory of distributed, collaborative production communities and networks is needed. It would be good to have a more systematic terminology for these various forms of trust-based collaboration, social production, distributed, networked, egalitarian communities, let alone a grand theory thereof. A theory should outline the archetypes, their characteristics, their strengths and weaknesses, the required circumstances, their viability, and much more.

Benkler's model apparently lacks predictive power and does not give any indications when peer production, other variants of trust-based collaboration, or some mixtures of trust-based and authority- or money-based forms of collaboration can be expected. Efficiency and effectiveness, the concepts that according to Benkler make peer production viable, are blind to the political subtleties of realpolitik, material interest, and urge for power, concepts that play an important role in security politics. Institutional economics has tremendous analytical power as the success of the school of the economics of information security has shown in recent years. Yet, Internet security policy has become a matter of national interest, of geopolitics, of International politics. It would therefore be worthwhile to approach the empirics of Internet security communities from the angle of established IR theories. Furthermore, it should be worthwhile to analyse the security communities using theories of epistemic communities.

8.4 The state of Internet security governance

In this remaining section, I'd like to contextualise my research within wider questions of Internet security. In a way, the empirical chapters of this study are less an inquiry into questions of public policy and contemporary Internet governance, and more a journey into the near past of what may eventually be characterised as the time when Internet security began to strongly and visibly overlap with questions of national security. Indeed, the public discourses in both cases had characteristics of traditional national-security debates: national states and their institutions as identified or assumed actors behind threats to what has been labelled as national security interest.

8.4.1 Technological cosmopolitanism vs. the information umbrella

The influx of national security pundits into Internet security discourses has fundamentally changed perspectives on Internet security. In 2009, mainstream discourse on Internet security governance circled around the idea of fire brigades or similar organisational arrangements. In a TED talk, Jonathan Zittrain stated that the Internet would be kept up and running by “random acts of kindness” by “unsung heroes”. To deal with incidents that undermine the Internet’s technical integrity, “random people appear from nowhere, put out the fire and leave without expecting payment or praise”, Zittrain added (Fildes 2009). At the Internet Governance Forum in Egypt later that year, Vint Cerf, TCP/IP standard co-author turned Google *Internet Evangelist*, picked up that idea of people putting out blazes and called for a global Internet police brigade (author's observation). The analogy of fire brigades raises a number of interesting, yet contradictory implications. The fascinating history of the organisation of fire response teams in the U.S. and its preceding colonial territories took a turn the ICT industry would not appreciate for incident response teams. Initially, the fire fighting business was in the hands of volunteer, yet elitist and politically very influential fire companies, which attracted the likes of Benjamin Franklin and George Washington (McChesney 1986, p. 73). Later, insurance companies started compensating these fire companies for mitigating fires at houses insured by them, which in turn attracted the unemployed, hangers-on, boxers, and thugs, resulting in “violence over the right to fight” fires (1986, pp. 77-78), which in turn lead to the rise of municipal monopolies, supported by public demand. With the emergence of Stuxnet as the latest, the idea of voluntary fire brigades as the protection against action against national cyber-warfare machineries appears to be glaringly disproportionate.

The situation today is an awkward co-existence between entirely different organisational approaches to Internet security governance: from the once “unsung heroes” whose not-so-random-at-all activities and communities have now finally been described and analysed in this study, over the slow rise of high-tech crime teams in police bureaucracies, and the build-up of offensive cyber-warfare capabilities, to the submission of the Internet’s giants to the status of mere agents in a fully-fledged surveillance regime driven by national intelligence agencies. The “shadow of hierarchies” has been turned into a brightly lit elephant in the room. It is the US government and its national security institutions that have the capacity to see what is happening on the Internet. Many of the security services provided by the Internet security community could arguably be replaced by services from NSA-fed national security circles in the medium-term. There are certainly reasons for the US government to continue to play a reduced role in the domain of infrastructural security and let the Internet security community continue to do much of the work as described in the empirical chapters of this study. Internet security has many facets and they can certainly be addressed by a variety of sometimes conflicting institutional approaches. But the revelations of the enormous capacity of national security institutions may alter the characteristic of the Internet security community: from “unsung heroes” who have voluntarily provided the global good of an up-and-running Internet to useful agents that may voluntarily provide public services as long as they fit into overall governmental security strategies.

Studying Internet security now requires that we revisit aspects of international relations. True, there are a number of aspects of Internet security in the areas of cybercrime and infrastructural security that do not raise national security concerns. But as the Internet has become a means, and an object of, vulnerabilities for national security interests, Internet security governance faces the same global governance and institutional issues that traditional security has.

Known for texts usually as dense and legible as the obfuscated JavaScript files of Google's (or Twitter's or Facebook's…) web applications, Immanuel Kant sardonically summed up the roots of the security problem of the human species: they are a "group of persons who can't do without peaceful togetherness and yet can't help but constantly loathe one another" (Kant 1796/2008, pp. 331; translation by author). The solution the saddler's son proposed in the late 18th century was to follow the guiding principle of the idea of a cosmopolitan society, in which citizens set up their rules and mutually ensure the following of these rules (1796/2008, p. 331). Apparently, Kant's approach to overcoming the security problem was to globalise civic rationality and reason, and thereby both globalise the civic ideals and set the precondition for a civil society (Richter 1992, pp. 37-55).

The Internet security community arguably produces security in a way that manages to bring together reason and globality. In the cases analysed in this study, the community re-established the Internet's functionality and thereby provided security as a global public good. The cases exemplified what could be described as technological cosmopolitanism. The point of reference, in security studies lingo: the security object for the Internet security community is not humanity in its direct sense, but the borderless, global Internet and its functionality as a technical artefact. The Conficker response showed a stunning global collaborative effort to contain the risks that were inherent in this tremendous botnet, mostly irrespective of conflicts between the countries of some participants. The Estonian cyberattacks have probably been the first incident in the history of nation states in which a Minister of Defence stated that his country was under attack and that a national security situation existed, and yet the military had no role whatsoever in defending the country from such an attack.

Students of international relations and political theory have long been haunted by the question of how to establish, create, and eventually protect the values of democratic, liberal societies in ways that do not endanger their civil pillars and that do not run afoul of their societies' humanistic values. In his book on the "Dissolution of world unity”, Richter argues that globality and reason have only been in accord with Kant’s ideas on cosmopolitanism. The actual characteristics of globalisation however have severely limited the possibilities to further reason on a global scale. As a consequence, theorists have come up with various models to further reason and civic values in an imperfect world, resulting in increasingly less ideationally ambitious models of global governance in order to preserve a chance of realizability (1992, pp. 242-252).

It is as if Internet security governance is moving in the opposite direction, from technological cosmopolitanism all the way to national, unilateral, even hegemonic models of Internet security. The global Internet security governance landscape looks very different than it did in 2009. That the state had begun to be massively involved in the field of cybersecurity was still new and noteworthy in 2010 (Mueller 2010, p. 179). Autocratic countries have long been criticised for their attempts to meddle with the technological systems that are part of the Internet within their national territories. With Stuxnet, and the revelation of its backers (US and Israel) and the surveillance practices of Western intelligence agencies it became public knowledge that any major Internet system or service is an object of and resource for national security policies, even in those countries with a yen for democratic values.

Before he actually began working on the design of the TCP/IP protocol, Kahn had written down a list of requirements which the protocol should fulfil. One has been summarised in a joint paper authored by the scientists that laid the technical foundations of the Internet: “There would be no global control at the operations level.” (Leiner et al., 2012) While TCP/IP is still largely the same, “global control” at that level is not inconceivable any more. The US has set a stunning precedent with its capability to access large chunks of the Internet’s raw traffic data and to the data of its domestic IT industry, which dominates the global IT and Internet industry. With monitoring capabilities as exposed by the Snowden leaks, the US appears to have created the “information umbrella”, that Joseph Nye and William Owens proposed as the information age successor to the “nuclear umbrella”, which served the US to both protect itself and its allies, while thereby increasing the latters' dependence within the security domain.

> “These capabilities [dominant situational knowledge] point to what might be called an information umbrella. Like extended nuclear deterrence, they could form the foundation for a mutually beneficial relationship. The United States would provide situational awareness, particularly regarding military matters of interest to other nations. Other nations, because they could share this information about an event or

> crisis, would be more inclined to work with the United States… Just as nuclear dominance was the key to coalition leadership in the old era, information dominance will be the key in the information age." (Nye & Owens, 1996, p. 27)

Regarding the necessity to reconfigure remote systems to deal with some Internet security issues, some countries have discussed and proposed to grant their law enforcement authorities the right to alter remote machines.[141]

8.4.2 The tragedy of national Internet securities

Even though the trajectory of current developments in security governance points toward increasingly national, if not unilateral or even hegemonic approaches to Internet security governance, the future is still unknown. Given the prevalence of realist and neoliberal thinking in foreign policy and international relations, it appears likely that traditional security institutions will try to further their say in Internet security governance. On the other hand, such an approach is likely to stir some public opposition given its detrimental effects on civil values. Furthermore, the Internet security community with its peculiar organisational rules, structures, and functions still plays a decisive role in those segments of Internet security that deal with botnets, cybercrime, and large-scale incidents. Eventually, any future model of Internet security will more or less resemble either of the five traditional models of security governance: cosmopolitanism, state-centric approaches, security regimes or communities, global civil society, or unilateralism (Loader & Walker, 2007). Today, all of these approaches coexist on different levels (subnational, national, international, transnational) in different domains of Internet security (cyberwarfare, cyberterrorism, cybercrime).

But all these IR-based considerations on security still do not take the networked approach into account. In my "Hierarchies in networks" article, I have argued that traditional security institutions, i.e., the NSAs, FBIs and militaries, might be able to achieve a *primus inter pares* position within an otherwise egalitarian system of mostly technical security experts (Schmidt 2014). While the security community and law enforcement both seek to address criminal usage of the Internet, there is a significant cultural clash between national security experts and those in the Internet security community. While the latter adheres to the idea of technological cosmo-

[141] For further discussions on the feasibility and potential characteristics of an Internet security hegemony compare (Schmidt 2014).

politanism, national security planners perceive the Internet as a patchwork of national fragments and a platform to achieve advantages in the international system.

As a consequence of the NSA revelations, Jeff Moss, initiator of the Defcon hacker conference series, asked the feds to stay away from the 2013 event (Franceschi-Bicchierai 2013). Given that at the same time the conference provided Keith B. Alexander, head of the NSA, an opportunity to talk about the surveillance practices of his agency in a keynote speech (Alexander 2013), it is doubtful that the symbolic request will result in a fundamental shift of attitudes among the US hacking community towards traditional security organisations. The budgets of national security institutions attract experts and the security industry alike. Some operational Internet security communities and hacker circles have close links to law enforcement. It is entirely feasible that these communities follow monetary and hierarchical incentives rather than intrinsic motivations such as supporting the cause of technological cosmopolitanism. In the Netherlands, the 2013 OHM hacking conference revealed a shift of attitudes, away from anti-establishment, subcultural hacking towards an ever closer collaboration with law enforcement and the security industry (Borchers 2013).[142] Making fun of stupid domestic corporations and governments by exposing their idiotic security design decisions is on its way out. The new cool is exposing cyber criminals and unveiling alleged or real Chinese/Russian/Iranian/Korean cyber attacks. Humanitarian hacking unmasks the involvement of authoritarian regimes in compromising ICT systems used by dissidents, which at the same time nurtures an anti-Chinese/Russian sentiment in public discourse. This also helps to legitimise the build-up of national offensive cyber-capabilities.

One could rate these developments as an indication of a maturation in Internet security policies. After all, societal and cultural innovation often begins in subcultures before it is adopted by incumbent institutions. However, the securitization of the Internet bears an element of tragedy for liberal societies. A situational awareness achieved through the ubiquitous collection of Internet-based data becomes a cogent necessity, when Internet-based data is regarded as a security risk or a strategic advantage. The tragedy of the security strategy — secretive mass surveillance to gain informational advantages against terrorists, adversaries, and untrustworthy allies — is that it undermines individual privacy, trust in ICT, and the digital foundations of the public sphere. Such a strategy upsets the public and their security institutions; it has the potential to undermine the values the strategy aims at

[142] An example of the critique from hackers with a rigorous ideological stance is groente 2013.

protecting in the first place.[143] Furthermore, it may stymie hopes that the Internet[144] could help to democratise autocratic countries.

Decisive steps have been taken by national security authorities without prior consulting of parliaments, let alone the public. Through this process, operational staff no longer sole technical authority over Internet components. Instead, authority has partly been transferred to national security organisations. In his review of Milton Mueller's *Ruling the Root* in *Salon* magazine, Andrew Leonard wrote, "as a narrative", the book was an "account of how the geeks eventually lost control of their creation" (Leonard 2002; Mueller 2002). In the responses described in this study, the geeks were still acting as the sole guardians of the Internet and jointly controlled the Internet's decisive components. Since then, however, the locus of authority and the coercive potentials have arguably changed. Administrators of decisive Internet systems have apparently lost their monopoly of control over ICT systems and can be overridden by third parties. Such remote take-overs are facilitated by a seemingly never-ending stream of new zero-day exploits and an ICT industry, that voluntarily or compelled acts as data delivery agent for its national security principals. The Internet security community does not act in the shadow of hierarchy, but in the shadow of hegemony. That being said, there still is a significant role for the Internet security community. In incident response, e.g., this community may indeed continue to play a decisive coordinative role, and contribute decentralised, denationalised, and distributed aspects of security.

The potentially counteracting effects of the national, as well as the hegemonic model of security production on civic values, are familiar to any student of security organisations. Loader and Walker summarise the delicate relation between security and means of security: "As monopoly holders of the means of legitimate physical and symbolic violence, modern states possess a built-in, paradoxical tendency to undermine the very liberties and security they are constituted to protect." (2007, p. 7) The negative effects of monopolistic security institutions on Internet security have been anticipated. Internet scholars have consequentially recommended security institutions with only moderate authority. For example, Mueller called for an approach that followed the idea of a "denationalized liberalism" (Mueller 2010, pp.

[143] Of course, there are likely some substantial gains: knowledge about adversaries and competing countries.

[144] "Once strong-arm regimes open the door to technology, they may find it difficult to return to a culture of bureaucratic secrecy, unscrupulous abuse of power, and unaccountability. (...) Authoritarian governments may not enter the information age with reform in mind, but it can be a welcome result." (Kalathil 2003)

268-271), to mitigated concerns about the need for a safeguard against an intrusive central authority. He proposed “to make Internet users and suppliers an autonomous global polity”, so that they can bring about the “authoritative interventions” that are “needed to secure basic rights against coercive attacks”. The “globally networked communities”, among which the Internet security community could be classed, would eventually establish new “forms of control” (2010, pp. 268-271). Deibert and Crete-Nishihata have similarly called for an Internet security polity that furthers “transparency, accountability, and mutual restraint” (2012, p. 274). “Distributed security” would “help mitigate unchecked and concentrated political power” by dividing authoritative control among multiple actors and thereby making it impossible for just one actor to assert control over the Internet (2012, pp. 272-273). Leaving questions of realizability aside, the following section contemplates a strengthened role for the Internet security community in the general Internet security architecture.

8.4.3 Reasons for the community's relatively decreased role

The current state of Internet security governance supports the hypothesis that the trust-based, rather egalitarian, global Internet security community is in relative decline. The communities themselves are still engaged in countering threats to the technical integrity of the Internet. But since 2005, when the first anti-botnet gatherings took place, the institutional architecture of Internet security has changed. The role of the community has been reduced by a number of trends. Governments and national authorities have become increasingly involved in Internet security affairs, military cyber-capacities have been built up, and more formal public-private partnerships have been arranged. In addition, large corporations act as *primus inter pares* within the community. To fully back these casual observations on the security community, it obviously requires detailed analyses of distinctive communities, of their development over time, and a fully-fledged systematic assessment of existing governance approaches. None of this can be done in the remainder of this study. Nevertheless, the following builds on the assumption of the decreased role of security community.

Before contemplating the future role of the Internet security community, the reasons for its neglect in previous policy and security architecture design decisions need to be evaluated. Lacking a thorough empirical analysis about this question, a number of potential reasons with varying likelihood are evaluated.

One possible reason could be reduced awareness about the community and the networked security architecture among policy-makers and the public. A variety of indicators back this argument. First, there have been very few articles on the role of the Internet community, and most of them have only been published within the last years (cf. section 2.4 for a summary of the previous literature on Internet-related security communities). Secondly, important areas in the field of Internet security governance still have to be researched. Among other blank spots, there is little information about the effectiveness, efficiency, and scalability of the community-networked approach to Internet security production. Furthermore, we lack comparisons between the more informal community approach and formal information sharing institutions such as ISACs. Thirdly, for many years, policy-makers would only rarely pay attention to conferences, workshops, and meetings in which issues regarding Internet security governance were discussed. Furthermore, the Internet security community has preferred to remain under the public radar and also struggled to be heard in policy-making circles. Fourthly, some important questions are left unanswered regarding the networked approach to Internet security governance. The sustainability of this approach, its scalability, its integration with state-driven approaches, as well as the quality and assuredness of its services need to be evaluated.

The second possible reason why the community has been ignored in many recent Internet security debates could be that the informal-expert-networks-approach is deemed inappropriate by decision makers. Again, various factors support this claim. Firstly, policy makers have opined that the self-governance approach by the private sector, which partly relies on the Internet security community, is insufficient. Consequentially, they have backed legislations that introduced more regulations into the Internet security domain. The main proponents for this argument about the insufficiency of existing security governance have been the rising number of attacks on ICT systems. These strikes, which are initiated by criminals, aggressive hacktivists, or possibly foreign governments, have inflicted substantial damages upon citizens, companies, and public authorities. As a second factor, the cases analysed in this study could be, at least superficially, regarded as an argument for increased intervention. After all, community members said after the Estonian incident that they were lucky. If the attackers had slightly modified their attacks, the Estonians would have registered higher damages. The Conficker attack exhausted the communities' human resources, making the response approach unsustainable in the long run. Thirdly, from a realist or neoliberal IR perspective, cyberspace has become a contested terrain, where a power vacuum offers states opportunities to reap advantages over competing or hostile countries. However, competing countries could likewise offset shortcomings in other domains of power by building up excellent cyber-capabilities. Therefore, the threats emanating from foreign countries cannot be mitigated by a volunteering community. They rather require big investments into specific, specialised professional security institutions. As a fourth factor, a rising number of professionals with a background in traditional security institutions hold positions in large ICT companies and policy boards on Internet security issues. Prior to this influx, technical perspective on security issues dominated. This is now supplemented by perspectives of traditional security institutions. As a fifth possible factor, the Internet security community could be seen as a kind of vanguard that explores the feasibility of different technical and organisational approaches to deal with various Internet security issues. In the long run, Internet security would, according to this argument, resemble traditional forms of security governance.

The third possible reason for the reduced role of the community in policy debates is political: networked, distributed security governance is not the preferred mode among decision makers and therefore not supported. Political and transnational networks maybe *en vogue* among scholars, but not among practitioners. Bureaucracies are known for developing their own life when it comes to defending or expanding their areas of responsibility. The turf wars among the different branches in the US security apparatus have been quite illustrative. Furthermore, the wish for safeguards against Twitter revolutions in times of economic uncertainty may exist in democratic countries, too. According to the former German chancellor Helmut Schmidt in November 2012, revolution was "in the air" and "unexpected things" could happen in the years to come. A fourth factor as to why politicians may not prefer a more significant role for communities is: strategy. US foreign policy pundits and strategy planners had already identified the value of information technology as a driver and means of US interest back in the 1990s. The US strategy regarding Internet security policy has two very different angles. A laissez-faire strategy has facilitated self-regulation and the rise of a decentralised, global technical community with increasing involvement of LEAs. The unilateral "information umbrella" strategy (Nye & Owens, 1996) has been implemented in the programmes described in the Snowden documents and relies on hierarchical institutions.[145] Finally, the fifth factor to limited support by policy-maker is that the Internet security community might lose its partly implicit, partly explicit backing by the various arms of the ICT industry. The largest firms among them might opt for the feudal security model, others might prefer more formalized models such as

[145] These two strategies are largely complementary, but they also have contradictory elements. International sharing in communities can be risky from a national security perspective, as contributors are only vetted by the community. The national security perspective prioritizes risk minimisation over low costs and effectiveness of security provisioning. Therefore, if the US or any other country alters its threat assessment, a different approach towards the Internet security community might follow.

traditional PPP relations or institutionalised intra-industry collaboration, as these models allow them to act more independently or promise increased revenues.[146]

These are the potential reasons why the security production and governance approach that is represented by the Internet security community might not have played a prominent role in policy debates in the last couple of years. While there might indeed be a ray of truth in many of the above arguments, they are still essentially hypotheses. There are a number of reasons to be sceptical about the long-term feasibility of the Internet security community as a decisive instrument against Internet insecurities. But a convincing, thorough argument against the feasibility of the decentralised, networked, non-or-partial-hierarchical approach to Internet security governance has yet to be written.

8.5 Opening security — a civilising role for the community

The trajectory of current developments points towards a takeover of Internet security governance by technology seigneurs, national security authorities and unilaterally acting governments. However, we are not there yet. What we have is a simultaneity of fundamentally different approaches to Internet security governance, including: self-governance, spontaneous ordering, peer governance, code ("code is law") and technical Internet architecture, markets, national regulation, international organisation, and transnational institutions (Schmidt 2009, pp. 11-16). In some areas of Internet security, such as traffic monitoring or offensive military capabilities, states have amassed excessive capabilities. In other areas, however, they are almost helpless to respond adequately to existing risks. Many police high-tech crime teams, for example, still only investigate criminal activities with damages exceeding a threshold of hundreds of thousands of Euros.[147] In contrast to national intelligence organisations, which have centralised repositories of data coming from all sorts of sources, investigations into botnets and other forms of cybercrime still rely on distributed knowledge and authority over distributed machines. This re-

[146] Feudal security and close state-industry collaboration are not contradictory, though, if one. In the historic analogy, princes, counts, and other members of the higher nobility acted relatively autonomous within their designated territory or fiefdom. Castles are the most apparent architectural means to protect noble power. At times however, during the Holy Roman Empire, the king or later the emperor called to arms for the Reich's cause. At the king's request, overlords were obliged to perform their fealty or more precisely their expeditionary duty (*Heerfahrtspflicht*) by joining the king's non-standing army for a military expedition (Krieger 1992, pp. 27-31).

[147] E.g., in Germany. Source: informal background talk with an ICT forensics expert in 2013.

maining section, therefore, questions how Internet security governance can be designed to adhere to the principles of distributed authority over systems, egalitarian network structure, and Internet security as a global common good. This section seeks to adumbrate ways to strengthen the role of the Internet security community and to civilise security production.

Ruminating upon the 'civilising' effect of the Internet security community, there are some implied normative or hypothesising assumptions. The Internet security community arguably is a normatively preferable approach to others for the following reasons. Firstly, it offers a somewhat more open and transparent, less exclusive and coercive approach to security governance than traditional security institutions. While the community is far from being entirely open and transparent, it at least allows its members to defect by walking and talking away without being threatened by the harsh sentences which members of traditional security institutions can face. Secondly, the Internet security community, with its trust-based coordination mechanism, is less prone to scheming than hierarchical organisational forms combined with a strong authoritative and centralised force. Thirdly, the diversity and distribution of actors in the Internet security community allows for better oversight and checks-and-balances than other security governance models on the global scale. Even on the national level, democratic oversight over security organisations is often delegated to non-transparent groups of national parliaments. A global authoritative security would amass dramatic degrees of authority for which no adequate checking-and-balancing institutions exist. A unilateral approach by a democratic state would however possibly result in the cybersecurity equivalent of taxation-without-representation.

Next to these normative assumptions, this section further builds on some of the possible reasons of the hypothesised decreasing relevance of the Internet security community. First, the reduced role of the Internet security community is the consequence of the increased role of traditional security institutions and the expansion of Internet security domains into espionage, warfare, and counter-terrorism. The absolute role of Internet security governance is arguably not reduced, but its relative role is. Furthermore, a key cause of the relatively poor role the community has played in policy discourses is the lack of knowledge about the Internet security community and the networked approach.

Based on these assumptions, this section outlines four ways by which the Internet security community and the networked model could be strengthened and enhanced: a wider public discourse on Internet security governance, an enhanced public perception of the community, several means to strengthen the community's effectiveness and output, and last but not least, avoiding and countering developments with centralising effects on security governance.

8.5.1 A wider security discourse

It can be argued that public debates on Internet security frequently follow a similar pattern: an attack happens, its costs and impact are debated, how it could have been avoided and what it means for the future. Often, such assessments are followed by calls for new institutions and governance approaches, often including increased responsibilities or even coercive authority for traditional security institutions or formal public-private partnership arrangements. The fact that an attack happens supports the assessment that the networked approach to security production has not been feasible.

After Estonia, there was only little discussion about if and how the existing security governance model could be adapted to meet similar attacks in the future. This omission is all the more surprising as the response efforts have by-and-large been quite effective in the cases analysed. The implicit post-attack conclusion appears to be that the then existing security production model has failed and therefore requires an organisational replacement. Hitherto, both security studies and Internet governance studies have ignored this discussion. Furthermore, the costs of Internet threats are still vague, frequently based on figures from actors with vested interests, resulting in reduced or, more often, inflated figures. According to probably the most authoritative study on the costs of Internet security and the threats to it so far, a substantial chunk of costs are indirect, i.e., lost efficiency gains or revenues by users and companies deciding not to use ICT system due to a lack of trust in them (Anderson et al., 2012). In addition, the varying costs of different security models are not thoroughly discussed in the literature on Internet security.

Given all the rhetoric on the Internet as a driver for innovation by policy makers, think tanks, and 'policy advisors,' it is somewhat surprising how little of that impetus is spent on innovating security governance. Information sharing and public-private sector collaboration is the most commonly proposed way forward, yet little thought is spent on the actual shaping of these public-private partnerships. As we now know, they span such different phenomena as the Internet security community with its trust-based coordination mechanisms, co-financed botnet mitigation centres, and secret surveillance programmes. The reality of public-private partnerships allows egalitarian, flat, and open structures, just as they also encompass secretive contracts between governments and private service providers.

The theoretical perspective chosen in this study, open source and peer production theories, point to another direction for governance models to take. Just as governments have discovered the utility of open source software for public information systems, they might also want to re-evaluate the peer production approach to achieve Internet security. Businesses, the financial industry and ISPs might want to re-evaluate their collaborative strategies with regard to Internet security governance, just as businesses have agreed before on sharing the code of commodity software.

8.5.2 Public perception of the community

The minor role the Internet security community plays as an object in public security debates is matched by the limited presence of members or representatives of these communities in the debates themselves. Therefore, the community, or some communities therein, might want to reevaluate their approach to the public and the policy process. With increased openness about itself, the community could educate the public and its representatives on its characteristics, services and functions for the production of Internet security.

Some in the community might argue that more transparency could undermine the functioning of these communities. Increased visibility would turn them into likely targets of now alerted 'bad guys.' They could as well argue that the community has no obligation to inform the public and thereby increase its visibility, as the community is an entirely voluntary institution with no responsibilities to anyone but itself. However, the role of the Internet security community has become publicly relevant, as all Internet services require a functioning Internet which itself depends on the activities of these communities. Therefore, it is understandable that the public and its representatives need to learn about the community and existing institutions that help produce Internet security. It is within the community and its task forces that decisions are made about the security of the Internet. The community therefore is no longer is just a collaborative space for experts from the industry plus some geeks and hobbyists, but has become a provider of vital services.

There are precedents of the technical community acting both in highly technical forums, for example standard setting, and in policy communities alike. The Internet Society is a prime example of this technological-political duality. Complementing its traditional technological tasks, it is increasingly playing the role of an advisor to regulators and law makers. By and large, most community members try to stay away from political issues. Nevertheless, some in the Internet security community apparently see a necessity to act not only as a coordinator in the background. In a presentation targeted primarily at potential candidates for the community, Barry Raveendran Greene described the operational security communities as "civil society" (Greene 2012b).[148] A malware researcher opined that the security community had "a responsibility towards civil society" and, as the context of his statement suggests, helps to expose governmental attempts to infiltrate citizens' machines.[149]

The idea of a civil society has gained popularity in social sciences with the increasing globality of decision-making and governance processes during the last decades. With the absence of a global people and citizenry that could serve as a direct or indirect legitimiser for global policies, other sources of legitimacy have been badly needed. This is where civil society organisations and other elements of global public act as the equivalent of national publics in national democratic countries. They help providing globalization and global governance structures with a whiff of democratic legitimacy.[150]

One could argue that the community and its affiliated organisations have a certain right to act, as the latter own a particular segment of the Internet and have direct and indirect contractual relations with users. Nevertheless, the label of 'civil society' would certainly help to mitigate the legitimacy issue that the community faces when it exerts behind the scenes authority on a global scale without democratically constructed legitimacy. The community's procedural legitimacy could therefore be strengthened by including representatives from non-technical civil society into the community. Beyond that, the community should evaluate its own civil society status and consider building up politically neutral policy advisory and watchdog functions.

8.5.3 Avoiding centralising effects

The approach of egalitarian, networked communities requires an avoidance of centralising trends in Internet security governance. Traditional security organisations have a history of mediocre (or worse) oversights, power abuses, and internal secrecy that make them a less-than-optimally suited instrument to take control of the cur-

[148] The conference presentation by Greene has been publicly available online for years despite being classified as "Operational Security Community or prospective members" only.

[149] Claudio Guarnieri, @botherder, Twitter, https://twitter.com/botherder/status/351106457829261312

[150] For a discussion of the role of civil society in security governance and Internet governance cf. Loader & Walker, 2007, p. 254; Hofmann 2005.

rent societies' public forums and communication infrastructure. Instead, the institutions to secure the Internet should remain as accessible, transparent and distributed as possible. However, there are a number of ways in which authoritative control over Internet components could be centralised. It is worth recalling the drivers for a decentralised, distributed approach of Internet security governance in the first place, in order to anticipate where change could be coming from. It is the transnational distributiveness of the ownership of systems involved in attack and defence, that have the knowledge and authority to understand and address the situation, thus making the stringent provision of Internet security a distributed, collaborative endeavour. Attempts to centralise Internet security governance therefore have to overcome one or many of these drivers of distributiveness.

As ownership of Internet components will continue to be distributed for the foreseeable future, the focus here is on the locus of authority and knowledge. This locus can be altered by restructuring the Internet security community, by technical innovation, and by legislative or normative changes.

At least in theory, the inner structure and governance principles of communities can be hierarchical. States and their traditional security institutions could aim to achieve more influential roles in the Internet security community and make the communities' internal rules and structures more amenable to their cause. But not only states could aim at an increased hierarchy of the communities. Larger companies can have both an interest and the ability to influence the agenda, internal practices, and principles of security communities.

While currently much of the information that is necessary in the course of the response to an incident rests with many geographically distributed actors, innovations in technology and business models can lead to more central, proprietary monitoring systems, which receive and compute data from numerous globally distributed sensors. Such systems would centralise the locus of knowledge and reduce the need for community support. However, centralised monitoring systems fed by distributed sensors do not necessarily result in centralised, proprietary knowledge. If the knowledge created by such systems is shared openly or within the security community (such as in the case of the Security Information Exchange created by the Internet Systems Consortium), it empowers the security community rather than creating new hierarchical elements.

Legislative changes could allow either public authorities or private actors to ignore the distributed ownership of Internet connected systems by granting them the right to remotely alter the configuration of any system that is involved in Internet-based attacks. Such legislation would centralise the locus of authority. Recent plans by the Dutch government exemplify this; police forces would retain the right to hack into any machine involved in DDoS attacks or a botnet. A similar approach would be a mandate for OS and software providers to force-update or reconfigure their installed base.

Conflicting norms structure the behaviour of members of the Internet security community. Introducing the nation state and its national security interests as points of reference in Internet security politics diverts the focus and loyalty of the community away from the goal of global infrastructural Internet security towards national interests and a parcelling of Internet security. As a consequence, Internet security becomes a patchwork rug of at times conflicting, at times complementary islands of national Internet securities. National Internet security requires a legally backed technical centralisation of the locus of authority. The normative idea of nationally responsible security hacking ensures the availability of required skills and marginalises progressive hacking scenes and their libertarian ideas.

Among the drawbacks of the centralised, hierarchical approach is that the introduction of hierarchical and monetary incentives might crowd out the intrinsic motivations of contributors. If this happens before hierarchical and commercial institutions are fully established, the overall security situation would worsen. As a second draw-back, the set-up and maintenance costs of a centralised, hierarchical governance system are likely to be substantially higher than those of the social/networked approach. And last, but surely not least, the potential externalities of hierarchical and central Internet security governance would be substantial. From the perspective of democratic governance, the societal risks of letting traditional security organisation manage the security of the Internet appear to be significant.

8.5.4 Strengthening its effectiveness

An important way to increase the relevance of the networked approach is to increase its capabilities and effectiveness. This would remove political doubts regarding its ability to handle the increasing variety and number of security incidents. The technical integrity of the Internet resides in the hands of the Internet security community and the organisations connected to it by default. The technical staff of individual organisations and their ICT service suppliers ensure the functionality of their networks and systems. One can conceive a number of ways in which security provisioning in this networked approach could be improved (Rendon Group 2011).

The community has a variety of options to increase its capabilities and effectiveness, without abandoning its key organisational feature, the trust-based coordination mechanism. A known weakness of social production is the delicate nature of intrinsic motivation. The latter can be hampered by a number of reasons: Being overburdened by too many issues at the same time (e.g., by three simultaneous Conficker-like incidents); being bored by too many tedious chores instead of challenging, inspiring puzzles; insufficient possibilities of ex-post indirect appropriation; crowding out intrinsic motivation by introducing monetary incentives or hierarchical compulsion. Consequentially, the community needs an environment that nurtures the motivations of its members in different situations.

Increasing the number of potential contributors to the efforts of the community can address some of these potential issues. Overall efficiency can be increased by providing a collaborative platform, and creating further community-open security tools. Redesigning work-packages and identifying tasks that require less trustful relationships can address the trust-scale problem. The motivation-chores problem can be addressed by decreasing the granularity of tasks and increasing the number of potential contributors. Ideally, the community is directly or indirectly connected to every ISP or at least to every ASN owner worldwide. Consequentially, the community needs to overcome its US/Western centrism and build deeper connections with security operators in other countries, even those with dubious political standards. The extended reach of the community should also include scientists in domains that are hardly represented in the community so far, e.g., mathematical cryptologists.

An apparent shortcoming and even vulnerability[151] of the networked-community model are occasional holes between communities. To respond to attacks with constantly changing combinations of attack technologies, scale of attacks, affected actors and geographies, each attack potentially requires a different set of security communities for the response endeavour. Therefore, gaps in the network of communities and the lack of interconnectedness between them can have detrimental effects. The Estonian case has exemplified the problems and importance of inter-communities collaboration.

The efficacy of social production elements within the overall response can, at least theoretically, be increased by creating new opportunities for openness and by in-

[151] "The vulnerability provides a description or the 'degree' relative to the connectivity of a community to: (1) other communities and (2) the network itself." (Rocco, Claudio, & Ramirez-Marquez, 2011, p. 1362)

stalling a standing infrastructure to support ad-hoc incident response activities. A community that defaults to secrecy loses one of the key efficacy drivers of peer production: that every person can chose the area and work package to which she wants to contribute. Defaulting to security was partly driven by efficiency considerations and lack of resources to evaluate the advantages of secrecy on a more granular level.

Granularly opening security data and collaborative communities might lead to the kind of innovations that are driving high-level Open Data policies in countries such as France (data.gouv.fr), the United Kingdom (data.gov.uk), the United States (data.gov) and the European Union (du Preez 2011; Huijboom & Van den Broek, 2011). This certainly increases the set-up costs for organising response in general, but could decrease marginal set-up costs of ad-hoc organisations for dedicated incidents and therefore decrease transaction costs for peer-producing initiatives.

While this loose coupling is partly a strength of the community, it is also a weakness. In some communities, members are not required to act upon the requests of their peers. Security is only a priority if it fits into the overall schedule and mood of a community member. However, some communities have already established contributions as mandatory in certain cases. To ensure reliable outcomes, the communities should consider committing themselves on certain quality standards and service levels.

In addition, the scalability problem of the security community is addressed by a variety of measures, among them a more nuanced approach to secrecy. The continuing reliance on an organisational hybrid which blends a social mode of production with a networked collaboration of technical companies, ensures an opportune provisioning of Internet security by avoiding the build-up of numerous Internet security bodies worldwide.

While the organisation-less, entirely informal approach has substantial benefits, the community should re-evaluate the experiences of other coproduction communities. In open/free software for example, communities often switched to a more formal approach to deal with problems of scalability, balancing contributors' different interests, and the need for more valid decisions.

The community is a mix of a coordinating body for independent responses by private organisations and a collaborative workbench to develop and implement joint solutions. Private organisations can contribute to the networked-communities approach by allowing a few persons of their ICT staff to collaborate and contribute to security communities. By attaching themselves to the security community, companies have better information about the overall security situation and have quick access to remediation forces in case of an attack. However, the rules of the communities and the split loyalty of their members can pose a cultural challenge to some organisations.

While the networked-communities approach stresses self-governance, bottom-up organisation, and regulatory independence, by no means does it imply that there is no role for the state. Its role is simply different, with public authorities acting as nodes in a more or less egalitarian network. In this model, nation states primarily contribute by allocating law enforcement resources to cover the Internet security production processes of identifying and sanctioning perpetrators. Most services that are necessary during the response endeavours — monitoring, problem analysis, mitigation — would still come from members of the security community and their affiliated organisations.

A number of weaknesses within the current implementation of the networked-communities model could be improved by the following measures. A major deficit of Internet security production has been the difficulty of identifying and sanctioning perpetrators of Internet-based crime. While international cooperation has apparently improved over the past years, law enforcement agencies worldwide still lack personnel, resources, and speedy transnational procedures to identify and arrest perpetrators. The build-up of these capabilities supplements the existing model as technical communities can hardly take over the sanctioning dimension of security. In addition, law enforcement can provide the intelligence the security community needs for its forensic analysis.

A fundamental weakness of the current institutional design arguably is a lack of knowledge among policy makers worldwide about the procedures and characteristics of the security community. Policy makers and the community need to find ways to deepen their mutual understanding. One way to achieve such mutual understanding is to establish a political-technical community consisting of members of the community and persons responsible for policy formulation in political bodies. A similar arrangement could be made with the civil society, which has also not been represented in existing Internet security governance institutions. The difficulties of cooperation between hierarchies and networks might be tackled by liaising the technical community with political authorities. However, it will require substantial changes in attitudes and culture in public authorities and in policy ranks. Networked governance already poses cultural challenges to bureaucracies, but security production by egalitarian networked-communities is networked governance on steroids.

As most states want to increase their overall level of Internet security, they could act as a facilitator of the community. The Rendon Group report on the Conficker response already concluded that the community's efficiency could be increased by institutionalising some aspects of the community, providing monetary resources, sharing premises, technical infrastructure, and administrative staff. To ensure that no single actor achieves a position to dominate, other collaborative communities have implemented certain limits that cut contributors rights at certain levels. Furthermore, states could support this new security production and governance approach by sponsoring research supporting the technical, organisational, and political underpinnings.

Bibliography

Aarelaid, H. (2007, September 19). Overview of recent incidents. [Presentation given at ENISA/CERT CC Workshop on Mitigation of Massive Cyberattacks 19th September in Porto, Portugal]. Retrieved November 8, 2010, from http://www.enisa.europa.eu/act/cert/events/files/ENISA_overview_of_recent_incidents_Aareleid.pdf

Aben, E. (2009). Conficker/Conflicker/Downadup as seen from the UCSD network telescope. Cooperative Association for Internet Data Analysis. Retrieved February 8, 2013, from http://www.caida.org/research/security/ms08-067/conficker.xml

Adams, S. (2009, January 25). Conficker Windows virus infects 15 million PCs. *The Telegraph*. Retrieved February 12, 2014, from http://www.telegraph.co.uk/technology/4338625/Conficker-Windows-virus-infects-15-million-PCs.html

Adler, P. S. (2001). Market, hierarchy, and trust: The knowledge economy and the future of capitalism. *Organization Science*, *12*(2), 215-234.

Adler, P. S., & Heckscher, C. (2006). Towards collaborative community. In P. S. Adler & C. Heckscher (Eds.), *The firm as a collaborative community: Reconstructing trust in the knowledge economy* (pp. 11-106). Oxford, New York: Oxford University Press.

Adomaitis, N. (2007, April 28). Estonia calm after Red Army site riots, Russia angry. *Reuters*. Retrieved October 15, 2012, from http://www.reuters.com/article/2007/04/28/us-estonia-russia-idUSL2873034620070428/

Alas, J. (2006, May 10). May 9 protestors call for removing Bronze Soldier statue. *Baltic Times*. Retrieved from http://www.baltictimes.com/news/articles/15345/

Alas, J. (2007, March 7). Reformists pull off surprise victory, consider dumping centrists. *Baltic Times*. Retrieved October 15, 2012, from http://www.baltictimes.com/news/articles/17358/

Alas, J. (2007, February 21). Soldier fails to sway elections. *Baltic Times*. Retrieved October 15, 2012, from http://www.baltictimes.com/news/articles/17358/

Alexander, K. (2013). Black Hat USA 2013 keynote. *YouTube* [Audiovisual Material]. Retrieved August 28, 2013, from http://www.youtube.com/watch?v=xvVIZ4OyGnQ_

Anderson, N. (2007). Massive DDoS attacks target Estonia; Russia accused. *Ars Technica*. Retrieved September 27, 2012, from http://arstechnica.com/security/2007/05/massive-ddos-attacks-target-estonia-russia-accused/

Anderson, N. (2013). *The Internet police: How crime went online, and the cops followed.* WW Norton & Company.

Anderson, R. (2001). Why information security is hard — an economic perspective. In *Computer security applications conference, December 10-14, 2001* (pp. 358 - 365). doi:10.1109/ACSAC.2001.991552

Anderson, R. (2002). Security in open versus closed systems – the dance of Boltzmann, Coase and Moore. *Open Source Software Economics*, 127-142.

Anderson, R., & Moore, T. (2006). The economics of information security. *Science*, *314*(5799), 610-613. Retrieved from http://www.sciencemag.org/cgi/content/abstract/314/5799/610

Anderson, R., & Moore, T. (2007). Information security economics—and beyond. *Lecture Notes in Computer Science*, *4622*, 68. Retrieved from http://www.springerlink.com/index/v7771248505010p6.pdf

Anderson, R., Barton, C., Böhme, R., Clayton, R., van Eeten, M. J. G., Levi, M., . . . Savage, S. (2012). Measuring the cost of cybercrime. [11th Annual Workshop on the Economics of Information Security WEIS 2012, Berlin, Germany, 25-26 June]. Retrieved from http://weis2012.econinfosec.org/papers/Anderson_WEIS2012.pdf

Anderson, R., Böhme, R., Clayton, R., & Moore, T. (2008). *Security economics and the internal market*. European Network and Information Security Agency (ENISA).

Anderson, R., Böhme, R., Clayton, R., & Moore, T. (2009). *Security economics and european policy*. CiteSeerX. doi:10.1.1.145.8444

Ansell, C. (2000). The networked polity: Regional development in Western Europe. *Governance: An International Journal of Policy and Administration*, *13*(3), 303-333.

Arbor Networks. (2009). Protecting IP services from the latest trends in botnet and DDoS attacks. [White Paper]. Retrieved from http://www.arbornetworks.com/index.php?option=com_docman&Itemid=974&task=doc_download&gid=309

Arnold, C. (2009, March 30). Russian group's claims reopen debate on Estonian cyberattacks. *Radio Free Europe / Radio Liberty*. Retrieved from http://www.rferl.org/articleprintview/1564694.html

Ashmore, W. C. (2009a). *Impact of alleged Russian cyber attacks* [School of Advanced Military Studies United States Army Command and General Staff College Fort Leavenworth, Kansas; similar version published in Baltic Security & Defence Review, Vol. 11, 2009, 4-40]. Retrieved from http://www.dtic.mil/cgi-bin/GetTRDoc?AD=ADA504991&Location=U2&doc=GetTRDoc.pdf

Ashmore, W. C. (2009b). Impact of alleged Russian cyber attacks. *Baltic Security & Defence Review*, *11*, 4-40. Retrieved from http://www.bdcol.ee/files/files/documents/Research/BSDR2009/1_%20Ashmore%20-%20Impact%20of%20Alleged%20Russian%20Cyber%20Attacks%20.pdf

Assaf, D. (2007). Government intervention in information infrastructure protection. In E. Goetz & S. Shenoi (Eds.), *IFIP International Federation for Information Processing: Critical infrastructure protection* (Vol. 253, pp. 29-39). Springer Boston. doi:10.1007/978-0-387-75462-8_3

Aviram, A. (2004). Network responses to network threats: The evolution into private cyber-security associations. *SSRN eLibrary*. doi:10.2139/ssrn.570342

Bachfeld, D. (2009, January 21). Wurm dringt in Systeme der britischen Armee ein. *Heise Security*. Retrieved February 13, 2014, from http://www.heise.de/security/meldung/Wurm-dringt-in-Systeme-der-britischen-Armee-ein-200621.html

Baldwin, D. A. (1997). The concept of security. *Review of International Studies*, *1*, 5-26. doi:10.1017/S0260210597000053

Barlow, J. P. (1996, February 9). A cyberspace independence declaration. [Usenet]. Davos. Retrieved March 13, 2004, from http://w2.eff.org/Censorship/Internet_censorship_bills/barlow_0296.declaration

Barr, S. (1958). Book review: World peace through world law, by Greenville Clark and Louis B. Sohn. *Lawyers Guild Review*, *18*(3), 138-139. Retrieved May, 2010, from http://heinonline.org/HOL/LandingPage?collection=journals&handle=hein.journals/guild18&div=41

Basurto, X., & Ostrom, E. (2008). *Beyond the tragedy of the commons*. Working Paper Series. Retrieved April 2, 2009, from http://ssrn.com/abstract=1304688

Bauer, J., & van Eeten, M. (2009). Cybersecurity: Stakeholder incentives, externalities, and policy options. *Telecommunications Policy*, 706-719. doi:10.1016/j.telpol.2009.09.001

Bauer, J., van Eeten, M., & Chattopadhyay, T. (2008). Financial aspects of network security: Malware and spam. [Presentation, ITU-T Study Group 3 Geneva, Switzerland, 2 April 2008].

Bauer, J. M. (2005). *Internet governance: Theory and first principles*. Preliminary draft, for purposes of discussion only.

Bauer, J., van Eeten, M., & Chattopadhyay, T. (2008). *Financial aspects of network security: Malware and spam*.

Bauer, J. M., & van Eeten, M. (2009). Cybersecurity: Stakeholder incentives, externalities, and policy options. *Telecommunications Policy*, 706-719. doi:10.1016/j.telpol.2009.09.001

Bauwens, M. (2005). The political economy of peer production. *1000 Days of Theory*, *1*. Retrieved August 12, 2009, from http://www.ctheory.net/articles.aspx?id=499

Bauwens, M. (2012a). Characteristics of peer production. *Keimform* [Web page]. Retrieved from http:// keimform.de/2012/characteristics-of-peer-production/

Bauwens, M. (2012b). Evolving towards a partner state in an ethical economy. In A. Botero, A. Gryf Paterson, & J. Saad-Sulonen (Eds.), *Towards peer production in public services: Cases from Finland* (pp. 34-49) [Aalto University publication series Crossover]. Helsinki: Aalto University. Retrieved July 15, 2012, from http://co-p2p.mlog.taik.fi/files/2012/06/p2p-public-services-finland-2012.pdf

Bauwens, M. (2012c). Some observations on peer governance. *P2P Foundation blog* [Web page]. Retrieved May 22, 2014, from http://blog.p2pfoundation.net/some-observations-on-peer-governance/2012/12/11

Bauwens, M. (2012d). The triune peer governance of the digital commons. In D. Bollier & S. Helfrich (Eds.), *The wealth of the commons—A world beyond market & state.* Levellers Press. Retrieved January 10, 2014, from http://wealthofthecommons.org/essay/triune-peer-governance-digital-commons

Bayley, D. H., & Shearing, C. D. (2001). *The new structure of policing: Description, conceptualization and research agenda*. US Dept. of Justice, Office of Justice Programs, National Institute of Justice.

Baylor, K., & Brown, C. (2006). *Killing botnets: A view from the trenches* [White Paper]. McAfee. Retrieved January 10, 2012, from http://www.cfoworld.co.uk/cmsdata/whitepapers/3498/mcafee_killing_botnets.pdf

Bendrath, R. (2003). The American cyber-angst and the real world — any link? In R. Latham (Ed.), *Bombs and bandwidth: The emerging relationship between IT and security* (pp. 49-73). New York: New Press.

Bendrath, R. (2007). Der gläserne Bürger und der vorsorgliche Staat: Zum Verhältnis von Überwachung und Sicherheit in der Informationsgesellschaft. *Kommunikation@gesellschaft*, *8*(7). Retrieved from *http://www.soz.uni-frankfurt.de/K.G/B7_2007_Bendrath.pdf*

Benkler, Y. (1998). The commons as a neglected factor of information policy. In *26th Annual Telecommunications Research Conference.*

Benkler, Y. (1999). Free as the air to common use: First amendment constraints on enclosure of the public domain. *New York University Law Review*, *74*, 354-446.

Benkler, Y. (2001). Coase's penguin, or, Linux and the nature of the firm. *Eprint arXiv:Cs/0109077* [Paper preseted at 29th TPRC Conference, 2001]. Retrieved March 12, 2009, from http://arxiv.org/abs/cs/0109077

Benkler, Y. (2002). Coase's penguin, or, Linux and the nature of the firm. *Yale Law Journal*, *112*(3), 369-446. Retrieved August 3, 2009, from http://www.yalelawjournal.org/images/pdfs/354.pdf

Benkler, Y. (2004a). Peer production of survivable critical infrastructures. [Paper presented at Telecommunications, Policy, and Research Conference]. Retrieved May 3, 2009, from http://web.si.umich.edu/tprc/papers/2004/340/Benkler%20Critical%20Infrastrcutures.pdf

Benkler, Y. (2004b). Sharing nicely: On shareable goods and the emergence of sharing as a modality of economic production. *Yale Law Journal*, *114*(2), 273-359.

Benkler, Y. (2006). *The wealth of networks: How social production transforms markets and freedom.* New Haven, London: Yale University Press.

Benkler, Y. (2013). Practical anarchism: Peer mutualism, market power, and the fallible state. *Politics & Society*, *41*(2), 213-251. doi:10.1177/0032329213483108

Berendson, R. (2007). Küberrünnakute taga seisavad profid. *Postimees*. Retrieved from http://www.tarbija24.ee/120507/esileht/siseuudised/258409.php

Bijl, J. (2012). Critical analysis of Microsoft operation B71. *Fox IT blog* [Web page]. Retrieved from http://blog.fox-it.com/2012/04/12/critical-analysis-of-microsoft-operation-b71/

Blackbird. (2012). SIRv12: The obstinacy of Conficker. *Microsoft Malware Protection Center - TechNet blogs* [Web page]. Retrieved February 13, 2014, from http://blogs.technet.com/b/mmpc/archive/2012/04/25/the-tenacity-of-conficker.aspx

Blank, S. (2008). Web war I: Is Europe's first information war a new kind of war? *Comparative Strategy*, *27*(3), 227-247. doi:10.1080/01495930802185312

Boas, T. (2000). Cyberpolitik: The information revolution and U.S. Foreign policy. [Web page] Carnegie Endowment for International Peace. Retrieved May 4, 2004, from http://www.ceip.org/files/events/cyberpolitik.asp?p=5&EventID=51

Bodin, J. (1955). *Six books of the commonwealth* [Abridged and translated with an introduction by MJ Tooley]. Oxford: Blackwell. (Original work published 1576)

Bollier, D. (1999). The power of openness — why citizens, education, government, and business should care about the coming revolution in open source code software. *Berkman Center for Internet & Society* [Research Publication No. 1999-02].

Borchers, D. (2013). Dispute over police presence at OHM hacker festival. *The H security*. Retrieved August 22, 2013, from http://www.h-online.com/security/news/item/Dispute-over-police-presence-at-OHM-hacker-festival-1849264.html

Boscovich, R. D. (2012). Microsoft and financial services industry leaders target cybercriminal operations from Zeus botnets. *The official Microsoft blog* [Web page]. Retrieved from http://blogs.technet.com/b/microsoft_blog/archive/2012/03/25/microsoft-and-financial-services-industry-leaders-target-cybercriminal-operations-from-zeus-botnets.asp

Bott, E. (2012, April 15). The malware numbers game: How many viruses are out there? *ZDNet*. Retrieved February 18, 2014, from http://www.zdnet.com/blog/bott/the-malware-numbers-game-how-many-viruses-are-out-there/4783

Bowden, M. (2010). The enemy within. *The Atlantic*. Retrieved May 16, 2010, from http://www.theatlantic.com/magazine/archive/2010/06/the-enemy-within/8098/1/

Bowden, M. (2011). *Worm — The first digital world war*. New York: Atlantic Monthly Press.

Bowling, B., & Newburn, T. (2006). Policing and national security. [First Draft – Not for citation without authors' permission; sorry for that, AS], presented at London-Columbia 'Police, Community and Rule of Law' workshop, London 16-17 March 2006. Retrieved April 24, 2009, from https://clearingatkings.com/content/1/c6/01/84/31/policingandnationalsecurity.pdf

Brabham, D. C. (2008a). Crowdsourcing as a model for problem solving: An introduction and cases. *Convergence*, *14*(1), 75. doi:10.1177/1354856507084420

Brabham, D. C. (2008b). Moving the crowd at iStockphoto: The composition of the crowd and motivations for participation in a crowdsourcing application. *First Monday*, *13*(6), 1-22. Retrieved from http://firstmonday.org/htbin/cgiwrap/bin/ojs/index.php/fm/article/viewArticle/2159/1969

Brin, S., & Page, L. (1998). The anatomy of a large-scale hypertextual web search engine. *Computer Networks and ISDN Systems*, *30*(1-7), 107-117.

Brown, I., & Marsden, C. T. (2007). Co-regulating Internet security: The London Action Plan. [Paper presented at Second Annual Symposium of GigaNet, Rio de Janeiro, Brazil]. Retrieved December 10, 2008, from http://giganet.igloogroups.org/download/publiclibr/papers/ianbrown

Bryden, A., & Caparini, M. (2007). Approaching the privatisation of security from a security governance perspective. In *Geneva Centre for the Democratic Control of Armed Forces (DCAF): Private actors and security governance* (pp. 3-19). Münster: Lit Verlag.

Bumgarner, J., & Borg, S. (2009). Overview by the US-CCU of the cyber campaign against Georgia in August of 2008. *[US-CCU Special Report]*. U.S. Cyber Consequences Unit.

Camp, L. J., & Friedman, A. (2005). Peer production of privacy and security information. [Ethical Surveillance, 8-10 June 2005, Austin TX]. Retrieved from http://www.ljean.com/files/nEighborhood.pdf

Canavan, J. (2005). The evolution of malicious IRC bots. In *Virus bulletin conference* (pp. 104-114).

Caparini, M. (2007). Applying a security governance perspective to the privatisation of security. In M. Caparini & A. Bryden (Eds.), *Geneva Centre for the Democratic Control of Armed Forces (DCAF): Private actors and security governance* (pp. 263-282). Münster: Lit Verlag.

Cazemier, J., Overbeek, P., & Peters, L. (2010). *Information security management with ITIL, version 3.* Zaltbommel: Van Haren Publishing.

Cazemier, J. A., Overbeek, P. L., & Peters, L. M. (2004). *Security management* (7th ed.). London: The Stationary Office. (Original work published 1999)

Cerny, P. (2007). Multi-Nodal politics: Toward a political process theory of globalization. [Paper prepared for presentation at the annual conference of the International Political Economy Society, Stanford University, 9-10 November 2007]. Retrieved May 19, 2009, from http://www.princeton.edu/~smeunier/Cerny.doc

Charney, S. (2005). Combating cybercrime: A public-private strategy in the digital environment. In *Workshop on measures to combat computer related crime. 11th United Nations Congress on crime prevention and criminal justice, Bangkok* (pp. 18-25).

Christensen, C. M., & Rosenbloom, R. S. (1995). Explaining the attacker's advantage: Technological paradigms, organizational dynamics, and the value network. *Research Policy*, *24*(2), 233-257. Retrieved from http://linkinghub.elsevier.com/retrieve/pii/004873339300764K

Chumer, M., & Turoff, M. (2006). Command and control (C2): Adapting the distributed military model for emergency response and emergency management. In Van de Walle & Turoff (Eds.), *Proceedings of the 3rd international ISCRAM conference.* Newark, New Jersey.

Clark, D. D. (1992). A cloudy crystal ball — visions of the future. In *Plenary presentation at 24th meeting of the Internet Engineering Task Force, Cambridge, Mass* (pp. 539-543). Retrieved from http://www.ietf.org/proceedings/24.pdf

Clark, G., & Sohn, L. B. (1958). *World peace through world law.* Harvard University Press.

Clarke, R. (2000). 'Information wants to be free ...'. *Roger Clarke's web-site* [Web page]. Retrieved February 24, 2014, from http://www.rogerclarke.com/II/IWtbF.html

Clinch, J. (2009). *ITIL V3 and information security.* Clinch Consulting. White Paper. Retrieved November 10, 2013, from http://www.best-management-practice.com/gempdf/itilv3_and_information_security_white_paper_may09.pdf

Clover, C. (2009, March 11). Kremlin-backed group behind Estonia cyber blitz. *Financial Times.* Retrieved from http://www.ft.com/cms/s/0/57536d5a-0ddc-11de-8ea3-0000779fd2ac.html?nlcick_check=1

Coase, R. H. (1937). The nature of the firm. *Economica*, *4*(16), 386-405. doi:10.1111/j.1468-0335.1937.tb00002.x

Colby, W. E. (1976). Intelligence secrecy and security in a free society. *International Security*, *1*(2), 3-14.

Constantinou, A. (2010). Open is the new closed: How the mobile industry uses open source to further commercial agendas. *Open Source Business Review.* Retrieved September 26, 2013, from http://timreview.ca/article/330

Cooper, M. (2006). From wifi to WIKIS and open source: The political economy of collaborative production in the digital information age. *Journal on Telecommunications & High Technology Law*, *5*, 125. Retrieved December 5, 2008, from http://cyberlaw.stanford.edu/system/files/From+Wifi+to+Wikis+and+Open+Source.pdf

Cooperative Association for Internet Data Analysis. (2012). The UCSD network telescope. [Web page] Retrieved February 8, 2014, from http://www.caida.org/projects/network_telescope/

Council of Europe. (2001). *Convention on cybercrime*. Retrieved March 13, 2009, from http://conventions.coe.int/Treaty/en/Treaties/Html/185.htm

Cowhey, P., & Mueller, M. (2009). Delegation, networks, and Internet governance. In M. Kahler (Ed.), *Networked politics: Agency, power, and governance* (pp. 173-193). [Print] Cornell: Cornell University Press.

Crawford, A. (2006). Networked governance and the post-regulatory state? Steering, rowing and anchoring the provision of policing and security. *Theoretical Criminology*, *10*(4), 449. doi:10.1177/1362480606068874

Crowston, & Howison. (2005). The social structure of free and open source software development. *First Monday*, *10*(2). Retrieved September 27, 2013, from http://firstmonday.org/article/view/1207/1127

Daase, C. (1993). Sicherheitspolitik und Vergesellschaftung. Ideen zur theoretischen Orientierung der Sicherheitspolitik. In S. Feske, C. Schmid, B. Moltmann, & C. Daase (Eds.), *Regionalisierung der Sicherheitspolitik – Tendenzen in den internationalen Beziehungen nach dem Ost-West-Konflikt* (pp. 39-64). Baden-Baden: Nomos.

Dafermos, G. (2012). Authority in peer production: The emergence of governance in the FreeBSD project. *Journal of Peer Production*, (1). Retrieved January 1, 2012, from http://peerproduction.net/issues/issue-1/peer-reviewed-papers/

Danckworth, T. G. (2007). *Estlands Außenpolitik nach dem Beitritt zur Europäischen Union: Handlungsoptionen eines Kleinstaates.* Dissertation zur Erlangung des akademischen Grades doctor philosophiae (Dr. phil.) vorgelegt der Philosophischen Fakultät der Technischen Universität Chemnitz. Retrieved from http://www.qucosa.de/fileadmin/data/qucosa/documents/5778/data/diss.pdf

David, P. A., & Shapiro, J. S. (2008). Community-based production of open-source software: What do we know about the developers who participate? *Information Economics and Policy*, *20*(4), 364-398. doi:10.1016/j.infoecopol.2008.10.001

Davis, J. (2007). Hackers take down the most wired country in Europe. *Wired Magazine*, *15*(9), 15-09.

Deibert, R., & Crete-Nishihata, M. (2012). The nationalization of global cyberspace: Cyber security strategies of non-democratic states. In *International Studies Association convention 2012, San Diego California, US.* Retrieved from http://isanet.ccit.arizona.edu/MyISA/Validated/ConferenceItemDetailBasic.aspx?ItemID=36259

Den Besten, M., Loubser, M., & Dalle, J. M. (2008). *Wikipedia as a distributed problem-solving network*. University of Oxford. Oxford Internet Institute DPSN Working Paper Series No. 13. Retrieved January 12, 2010, from http://ssrn.com/abstract=1302898

Deshpande, & Riehle, D. (2008). The total growth of open source. In *Proceedings of the fourth conference on open source systems* (pp. 197–209). Springer Verlag. Retrieved January 28, 2010, from http://dirkriehle.com/wp-content/uploads/2008/03/oss-2008-total-growth-final-web.pdf

Dilling, O., Herberg, M., & Winter, G. (2008). *Responsible business : Self-governance and law in transnational economic transactions* . Oxford, Portland, Or.: Hart Publisher.

Dobbs, M. & Davies, A. (Screenwriters) & Seed, P. (Director). (2004). Episode 3. In R. Riddington (Producer). *To Play the King.* Part II of the House of Cards trilogy. [TV series episode]. England: BBC Worldwide Ltd. (Original work published 1993)

Dourado, E. (2012). *Internet security without law: How service providers create order online*. Mercatus Center, George Mason University. Working Paper No 12-19. Retrieved from http://mercatus.org/sites/default/files/ISP_Dourado_WP1219.pdf

Dresing, T., & Pehl, T. (2011). *Praxisbuch Transkription. Regelsysteme, Software und praktische Anleitungen für qualitative ForscherInnen* (3. ed.). Marburg. Retrieved from www.audiotranskription.de/praxisbuch

Drezner, D. W. (2004). The global governance of the Internet: Bringing the state back in. *Political Science Quarterly*, *119*(3), 477-498. Retrieved from http://www.ingentaconnect.com/content/taps/psq/2004/00000119/00000003/art00004

Drezner, D. W. (2007). *All politics is global — explaining international regulatory regimes.* Princeton, NJ: Princeton University Press.

Duffy Marsan, C. (2007, August 22). How close is world war 3.0? Examining the reality of cyberwar in wake of Estonian attacks. *Network World.* Retrieved from http://www.networkworld.com/news/2007/082207-cyberwar.html

Dunham, K., & Melnick, J. (2008). *Malicious bots: An inside look into the cyber-criminal underground of the internet.* Auerbach Pub.

Dunn Cavelty, M. (2007). Cyber-Terror – looming threat or phantom menace? The framing of the US cyber-threat debate. *Jounal of Information Technology & Politics*, *4*(1), 19-36. doi:10.1300/J516v04n01_03

Dunn Cavelty, M. (2013). From cyber-bombs to political fallout: Threat representations with an impact in the cyber-security discourse. *International Studies Review*, *15*(1), 105-122. doi:10.1111/misr.12023

Dunn Cavelty, M., & Suter, M. (2009). Public–private partnerships are no silver bullet: An expanded governance model for critical infrastructure protection. *International Journal of Critical Infrastructure Protection*, *2*(4), 179-187. doi:10.1016/j.ijcip.2009.08.006

Dupont, B. (2004). Security in the age of networks. *Policing and Society*, *14*(1), 76-91.

Dupont, B., Grabosky, P., & Shearing, C. (2003). The governance of security in weak and failing states. *Criminology and Criminal Justice*, *3*(4), 331.

du Preez, D. (2011, December 12). European Commission launches open data strategy. *computing.co.uk*. Retrieved December 17, 2011, from http://www.computing.co.uk/ctg/news/2131718/european-commission-launches-strategy-europe

Dutton, W. H. (2008). The wisdom of collaborative network organizations: Capturing the value of networked individuals. *Prometheus*, *26*(3), 211-230. doi:10.1080/08109020802270182

Dürrenmatt, F. (1994). *The physicists* [Die Physiker] (reedition). New York: Grove Press. (Original work published 1962)

Dye, R. A. (1985). Disclosure of nonproprietary information. *Journal of Accounting Research*, *23*(1), 123-145. doi:10.2307/2490910

Eckert, C. (2001). *IT-Sicherheit. Konzept - Verfahren - Protokolle.* München: Oldenbourg.

Egele, M., Scholte, T., Kirda, E., & Kruegel, C. (2008). A survey on automated dynamic malware-analysis techniques and tools. *ACM Comput. Surv.*, *44*(2), 6:1-6:42. doi:10.1145/2089125.2089126

Egele, M., Scholte, T., Kirda, E., & Kruegel, C. (2012). A survey on automated dynamic malware-analysis techniques and tools. *ACM Computing Surveys (CSUR)*, *44*(2), 6.

Egelman, S., Herley, C., & van Oorschot, P. C. (2013). Markets for zero-day exploits: Ethics and implications. In *Proceedings of the 2013 workshop on new security paradigms* (pp. 41-46). doi:10.1145/2535813.2535818

Eilstrup-Sangiovanni, M. (2007). *Varieties of cooperation: Government networks in international security*. Florence: European University Institute, Robert Schuman Centre for Advanced Studies. EUI Working Papers RSCAS 2007/24. Retrieved April 20, 2009, from http://cadmus.iue.it/dspace/handle/1814/7503

Eilstrup-Sangiovanni, M. E. (2005). Transnational networks and new security threats. *Cambridge Review of International Affairs*, *18*(1), 7-13. doi:10.1080/09557570500059498

Elgan. (2010). How Google is 'closed', just like Apple. *Datamation* [Web page]. Retrieved September 26, 2013, from http://www.datamation.com/features/article.php/3910226/How-Google-Is-Closed-Just-Like-Apple.htm

Estonia Hit by Moscow Cyber War. (2007, May). *BBC News*. Retrieved from http://news.bbc.co.uk/2/hi/europe/6665145.stm

Estonian DDoS - a Final Analysis. (2007, May 31). *The H Security*. Retrieved from http://www.h-online.com/security/news/item/Estonian-DDoS-a-final-analysis-732971.html

Etzioni, A. (2004). *From empire to community.* Palgrave Macmillan.

Etzioni, A. (2005). Genocide prevention in the new global architecture. *The British Journal of Politics & International Relations*, *7*(4), 469-484. doi:10.1111/j.1467-856X.2005.00201.x

European Commission. (2013). Proposal for a directive of the European Parliament and of the Council concerning measures to ensure a high common level of network and information security across the union. [COM(2013) 48 final]. Brussels. Retrieved May 1, 2013, from http://ec.europa.eu/information_society/newsroom/cf/dae/document.cfm?doc_id=1666

European Commission - DG Home Affairs. (2012). What we do - EU agency for large-scale IT systems. [Web page] Retrieved October 17, 2012, from http://ec.europa.eu/dgs/home-affairs/what-we-do/policies/borders-and-visas/agency/index_en.htm

Evron, G. (2008, May 20). [NANOG] an account of the Estonian Internet war. [ge@linuxbox.org, email sent to the NANOG mailing-list]. Retrieved January, 2011, from http://mailman.nanog.org/pipermail/nanog/2008-May/000676.html

Evron, G. (2008b). Battling botnets and online mobs. Estonia's defence efforts during the Internet war. *Georgetown Journal of International Affairs*, *9*(1), 121-126.

Evron, G. (2009, May 17). Authoritatively, who was behind the Estonian attacks? *Dark Reading.* Retrieved from http://www.darkreading.com/security/news/227700882/authoritatively-who-was-behind-the-estonian-attacks.html

Evron, G., & Aarelaid, H. (2008, January 17). Estonia: Information warfare and lessons learned. [Presentation given at the Workshop on Learning from large scale attacks on the Internet - Policy Implications]. Retrieved from http://ec.europa.eu/information_society/policy/nis/docs/largescaleattacksdocs/s5_gadi_evron.pdf

Feeny, D., Hanna, S., & McEvoy, A. F. (1996). Questioning the assumptions of the "tragedy of the commons" model of fisheries. *Land Economics*, *72*(2), 187-205. doi:10.2307/3146965

Felício, T. (2007). Multilevel security governance: Reinventing multilateralism through multiregionalism. *Human Security Journal*, *5*(Winter).

Feller, J., Finnegan, P., Fitzgerald, B., & Hayes, J. (2008). From peer production to productization: A study of socially enabled business exchanges in open source service networks. *Information Systems Research*, *19*(4), 475.

Ferguson, R. (2011). Conficker, Duqu, Stuxnet, aliens, confuxnet! *Trend Micro blog* [Web page]. Retrieved February 13, 2014, from http://countermeasures.trendmicro.eu/conficker-duqu-stuxnet-aliens-confuxnet/

Ferguson, R. (2012). Don't be dumb, keep schtumm! *CounterMeasures - Trend Micro blog* [Web page]. Retrieved January 10, 2013, from http://countermeasures.trendmicro.eu/dont-be-dumb-keep-schtumm/

Fildes, J. (2009). Unsung heroes save net from chaos. *BBC News.* Retrieved from http://news.bbc.co.uk/2/hi/technology/8163190.stm

Finkle, J. (2011, December 2). Insight: Did Conficker help sabotage Iran's nuke program? *Reuters Canada*. Retrieved February 13, 2014, from http://ca.reuters.com/article/topNews/idCATRE7B10AP20111202?sp=true

Finn, P. (2007). Cyber assaults on Estonia typify a new battle tactic. *Washington Post.* Retrieved from http://www.washingtonpost.com/wp-dyn.content/article/2007/05/18/AR2007051802122.html

Fohler, S. (2003). *Techniktheorien. Der Platz der Dinge in der Welt der Menschen.* München: Fink.

Foley, M. J. (2007, July 26). Microsoft: Expect Windows installed base to hit 1 billion by mid-2008. *ZDNet.* Retrieved February 12, 2014, from http://www.zdnet.com/blog/microsoft/microsoft-expect-windows-installed-base-to-hit-1-billion-by-mid-2008/596

Fossi, M., & others. (2008). *Symantec report on the underground economy: July 2007 to June 2008*.

Franceschi-Bicchierai, L. (2013). DefCon hacker conference to feds: 'We need some time apart'. *Mashable* [Web page]. Retrieved July 14, 2013, from http://mashable.com/2013/07/11/hacker-conference-defcon-no-feds/

Frankfurter Allgemeine Zeitung. (2007, June 18). Estland im Visier: Ist ein Internetangriff der Ernstfall? Retrieved November 4, 2010, from http://www.faz.net/s/RubDDBDABB9457A437BAA85A49C26FB23A0/Doc~E7CCF88CEFB6F467BB8D75A400C07B959~ATpl~Ecommon~Scontent.html

Frankie Goes to Hollywood. (1984). The power of love. On *Welcome to the pleasuredome* [LP]. Johnson, H., O'Toole, M., Rutherford, P., Nash, B., Gill, P., Richards, A., & Jardim, L. London: ZTT Records.

F-Secure. (2008). Trojan-Spy:W32/gimmiv.A. *F-Secure.Com* [Web page]. Retrieved February 11, 2014, from http://www.f-secure.com/v-descs/trojan-spy_w32_gimmiv_a.shtml

Galloway, A. R. (2004). *Protocol – how control exists after decentralization.* The MIT Press.

Gaycken, S. (2009, May). Die Zukunft des Krieges - strategische Konzepte und strukturelle Probleme des Cyberwarfare. [Paper presented at CCC's SIGINT conference, May 2009, Cologne]. Retrieved May 30, 2009, from http://events.ccc.de/sigint/2009/Fahrplan/attachments/1285_Die%20Zukunft%20des%20Krieges.pdf

German Instrument of Surrender. (2012). *Wikipedia: The free encyclopedia*. Retrieved October 26, 2012, from http://en.wikipedia.org/w/index.php?title=German_Instrument_of_Surrender&oldid=519649129

Giles, J. (2009, June 12). The inside story of the Conficker worm. *NewScientist.* Retrieved June 17, 2010, from http://www.newscientist.com/article/mg20227121.500-the-inside-story-of-the-conficker-worm.html

Goldsmith, J. L., & Wu, T. (2006). *Who controls the Internet? Illusions of a borderless world.* New York: Oxford University Press.

Gomez, W. (Transcript author). (2012). Building a secure cyber future: Attacks on Estonia, five years on. [Transcript of the ACUS workshop on May 23, 20012, Washington D.C]. The Atlantic Council of the United States. Retrieved August 24, 2012, from http://www.acus.org/print/70435

Goncharov, M. (2012). Russian underground 101. *Trend Micro incorporated research paper*. Retrieved November 4, 2011, from http://www.trendmicro.com/cloud-content/us/pdfs/security-intelligence/white-papers/wp-russian-underground-101.pdf

Gostev, A. (2009). The neverending story. *Securelist* [Web page]. Retrieved February 10, 2014, from http://www.securelist.com/en/weblog?weblogid=208187654

Gotlieb, R. (2011, March 14). Cybercop fights organized Internet crime. *Pacific Standard.* Retrieved from http://www.psmag.com/legal-affairs/cybercop-fights-organized-internet-crime-27897/

Gowder, P. (2005). Secrecy as mystification of power: Meaning and ethics in the security state. *I/S - A Journal of Law and Policy for the Information Society*, *2*, 1-25. Retrieved from http://www.is-journal.org/V02I01/2ISJLP001.pdf

Grant, R. (2007). Victory in cyberspace. [Special Report]. Air Force Association. Retrieved from http://www.afa.org/media/reports/victorycyberspace.pdf

Greenberg, A. (2009, May 29). What Obama's cyberplan means for business. *Forbes.* Retrieved February 13, 2014, from http://www.forbes.com/2009/05/29/cybersecurity-obama-business-technology-security-cybersecurity.html

Greene, B. R. (2003, August 16). NSP-SEC—peers working together to battle attacks to the net [Presentation]. Retrieved March 25, 2013, from ftp://ftp-eng.cisco.com/cons/isp/security/NSP-SEC/NSP-SEC-v1-2.pdf

Greene, B. R. (2012, February). The new Internet 'civic society'— OPSEC communities. [Presentation given at the Apricot 2012 conference, February 21-March 1, New Delhi]. Retrieved March 9, 2012, from http://www.apricot2012.net/__data/assets/pdf_file/0005/45293/005-Operational-Security-Community-2012-02-04.pdf

Greene, B. R. (2012, February). SP security primer 101—peers working together to battle attacks to the net! [Presentation given at the NANOG 54 meeting, February 5-7, 2012, San Diego]. Retrieved March 25, 2013, from http://www.nanog.org/meetings/nanog54/presentations/Sunday/Greene.pdf

Groenewegen, J. (2005, May 27). Designing markets in infrastructures: From blueprint to learning. [Professoral inaugurual speech, Faculty of Technology, Policy and Management]. Delft: Delft University of Technology.

groente. (2013). What's wrong with the kids these days? On the moral decay of the dutch hacker scene. *PUSCII blog* [Web page]. Retrieved May 28, 2014, from http://www.puscii.nl/blog/content/whats-wrong-kids-these-days

Gruszczak, A. (2008). Networked security governance: Reflections on the EU's counterterrorism approach. *Journal of Global Change and Governance*, *1*(3).

Habermas, J. (1990). *Strukturwandel der Öffentlichkeit - Untersuchungen zu einer Kategorie der bürgerlichen Gesellschaft.* Frankfurt am Main: Suhrkamp. (Original work published 1962)

Hafner-Burton, E. M., Kahler, M., & Montgomery, A. H. (2009). Network analysis for international relations. *International Organization*, *63*(3), 559-592. doi:10.10170/S0020818309090195

Hamburger, E. (2013, March 26). Apple's broken promise: Why doesn't iCloud 'just work'? *The Verge.* Retrieved March 11, 2014, from http://www.theverge.com/2013/3/26/4148628/why-doesnt-icloud-just-work

Hansabank Group. (2008). Annual report of Hansabank Group 2007. Retrieved January 4, 2011, from https://www.swedbank.ee/static/pdf/about/finance/reports/info_annual-report-2007_eng.pdf

Hansapanka Tabas Küberrünne. (2007, May 10). *Postimees*. Retrieved from http://www.tarbija24.ee/180507/esileht/majandus/259920.php

Hardin, G. (1968). Tragedy of the commons. *Science (New York, N.Y.)*, (162), 1243-1248.

Harrison, M. (2003). Why secrets? The uses of secrecy in Stalin's command economy. *University of Warwick. PERSA Working Paper*, *34*.

Hawkins, J., & Bates, T. (1996). Guidelines for creation, selection, and registration of an autonomous system (AS). IETF. RFC 1930. Retrieved February 10, 2013, from http://tools.ietf.org/html/rfc1930

Hänggi, H. (2003). Making sense of security sector governance. *Challenges of Security Sector Governance. Münster: Lit Verlag*, 3-22. Retrieved from https://intra.css.ethz.ch/mena/pub_haenggi_makingsense.pdf

Healey, J. (2012). Beyond attribution: Seeking national responsibility for cyber attacks. *Atlantic Council IssueBrief.* Retrieved April 3, 2012, from http://www.acus.org/publication/beyond-attribution-seeking-national-responsibility-cyberspace

Hemingway, E. (1936). On the blue water: A gulf stream letter. *Esquire, August*, *5*(4), 184-185.

Here We Go Again. (2007, April 4). *Baltic Times*. Retrieved from http://www.baltictimes.com/news/articles/17635/

Herzog, S. (2011). Revisiting the Estonian cyber attacks: Digital threats and multinational responses. *Journal of Strategic Security*, *4*(2), 4.

Hess, C., & Ostrom, E. (2003). Ideas, artifacts, and facilities: Information as a common-pool resource. *Law and Contemporary Problems*, 111-145.

Hoepman, J. H., & Jacobs, B. (2007). Increased security through open source. *Communications of the ACM*, *50*(1), 83.

Hofmann, J. (2005). (Trans-)formations of civil society in global governance contexts – two case studies on the problem of self-organization. In Schuppert (Ed.), *Schriften zur Governance-Forschung: Vol. 2. Global governance and the role of non-state actors.* Baden-Baden: Nomos-Verlag.

Hogben, G., Plohmann, D., Gerhards-Padilla, E., & Leder, F. (2011). *Botnets: Detection, measurement, disinfection & defence*. European Network and Information Security Agency (ENISA). Retrieved March 10, 2011, from http://www.enisa.europa.eu/act/res/botnets/botnets-measurement-detection-disinfection-and-defence/at_download/fullReport

Howe, J. (2006). The rise of crowdsourcing. *Wired Magazine*, *14*(6), 1-4. Retrieved from http://www.wired.com/wired/archive/14.06/crowds_pr.html

Huijboom, N., & Van den Broek, T. (2011). Open data: An international comparison of strategies. *European Journal of ePractice*, *12*(March/April 2011). Retrieved from http://www.epractice.eu/en/document/5290090

Hunderte Bundeswehr-Rechner von Conficker befallen. (2009, February 14). *Heise Security.* Retrieved February 13, 2014, from http://www.heise.de/security/meldung/Hunderte-Bundeswehr-Rechner-von-Conficker-befallen-195953.html

Hunfeld, F. (2008, December 19). Steuerfahndung Frankfurt - eiskalt abserviert. *Stern.* Retrieved March 16, 2009, from www.stern.de/politik/deutschland/649420.html

Hutter, B. M., & others. (2006). *The role of non-state actors in regulation.* Centre for Analysis of Risk and Regulation.

Huysmans, J. (2006). Agency and the politics of protection. Implications for security studies. In J. Huysmans, A. Dobson, & R. Prokhovnik (Eds.), *The politics of protection: Sites of insecurity and political agency.* New York, NY : Routledge.

Hyppönen, M. (2007, March 11). Allaple virus author sentenced. *F-Secure Weblog.* Retrieved February 23, 2014, from http://www.f-secure.com/weblog/archives/00001907.html

Hyppönen, M. (2007b). Unrest in Estonia. *F-Secure weblog* [Web page]. Retrieved from http://www.f-secure.com/weblog/archives/00001181.html

Hyppönen, M. (2007, April 30). Update on the Estonian DDoS attacks. *F-Secure Weblog*. Retrieved from http://f-secure.com/weblog/archives/00001183.html

Hyppönen, M. (2009, July 26). The Conficker mystery. [Black Hat Technical Security Conference USA 2009]. Retrieved February 14, 2014, from https://www.blackhat.com/presentations/bh-usa-09/HYPPONEN/BHUSA09-Hypponen-ConfickerMystery-PAPER.pdf

ICANN. (2009). The end of domain tasting, AGP deletes decrease 99.7%. *ICANN website* [Web page]. Retrieved February 18, 2014, from http://www.icann.org/en/news/announcements/announcement-12aug09-en.htm

IDC. (2013). PC shipments post the steepest decline ever in a single quarter, according to IDC. [Press release] [Web page]. Retrieved February 12, 2014, from http://www.idc.com/getdoc.jsp?containerId=prUS24065413

International Telecommunications Union. (2005). A comparative analysis of cybersecurity initiatives worldwide. [WSIS Thematic Meeting on Cybersecurity, Geneva, Document CYB/05] [Web page]. Retrieved April 23, 2009, from http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.115.1692&rep=rep1&type=pdf

Ioannidis, C., Pym, D., & Williams, J. (2009). Investments and trade-offs in the economics of information security. *Financial Cryptography and Data Security*, *5628*, 148-166.

Irion, K. (2012). The governance of network and information security in the european union: The European Public-Private Partnership for Resilience (EP3R). [27th European Communications

Policy Research Conference (EuroCPR), Policies For The Future Internet, 25-27 March 2012, Ghent, Belgium]. Retrieved from http://papers.ssrn.com/sol3/papers.cfm?abstract_id=2075916

Johnson, B. (2011, February 10). Nokia crisis highlights internal struggle. *BBC*. Retrieved February 10, 2011, from http://www.bbc.com/news/technology-12414595

Johnson, D. R., Crawford, S. P., & Palfrey Jr, J. G. (2004). The accountable Internet: Peer production of Internet governance. *Virginia Journal of Law & Technology*, *9*(9). Retrieved from http://ssrn.com/abstract=529022

Jung, F.J. (2009). The "networked security" concept – stocktaking and perspectives. *European Security and Defence*, (1), 7-12. Retrieved June 28, 2009, from http://www.europeansecurityanddefence.info/Ausgaben/2009/01_2009/01_Jung/ESD_0109_Jung.pdf

Jung, C. G. (2001). *Modern man in search of a soul.* Milton Park, Abingdon: Routledge. (Original work published 1933)

Kaeo, M. (2007). Cyber attacks on Estonia: Short synopsis. Retrieved from http://doubleshotsecurity.com/pdf/NANOG-eesti.pdf

Kahler, M. (2004, October). Global governance redefined. [Paper presented at The Conference on Globalization, the State, and Society, Washington University School of Law, St. Louis, 13-14 November 2003 (Revised October 2004)]. (Original work published November, 2003) Retrieved December 12, 2008, from http://irpshome.ucsd.edu/faculty/mkahler/GlobGov_10.04.doc

Kahler, M. (2009a). *Networked politics: Agency, power, and governance.* Cornell Studies in Political Economy. Cornell University Press.

Kahler, M. (2009b). Networked politics: Agency, power, and governance. In M. Kahler (Ed.), *Networked politics: Agency, power, and governance* (pp. 1-21). Cornell: Cornell University Press.

Kalathil, S. (2003). Dot com for dictators. *Carnegie Endowment for International Peace* [Printed originally in Foreign Policy Magazine] [Web page]. Retrieved June 9, 2004, from http://www.ceip.org/files/Publications/dotcomdictators.asp

Kant, I. (2008). *Anthropologie in pragmatischer Hinsicht* (Akademieausgabe von Immanuel Kants Gesammelten Werken). Korpora.org. (Original work published 1796) Retrieved from http://www.korpora.org/Kant/verzeichnisse-gesamt.html

Kaplan, D. (2008). Botnet experts meet as threat grows for corporations. *SC magazine* [Web page]. Retrieved February 10, 2014, from http://www.scmagazine.com/botnet-experts-meet-as-threat-grows-for-corporations/article/119773/

Kash, W. (2008). Lessons from the cyberattack on Estonia. Interview with Lauri Almann, Estonia's permanent Undersecretary of Defence. *Government computer news*. Retrieved from http://gcn.com/Articles/2008/06/13/Lauri-Almann--Lessons-from-the-cyberattacks-on-Estonia.aspx?p=1

Kaufmann, F. X. (1973). *Sicherheit als soziologisches und sozialpolitisches Problem: Untersuchungen zu einer Wertidee hochdifferenzierter Gesellschaften* (2., umgearb. Aufl.). Stuttgart: Ferdinand Enke Verlag.

Kazman, R., & Chen, H. M. (2009). The metropolis model a new logic for development of crowdsourced systems. *Communications of the ACM*, *52*(7), 76-84.

Keizer, G. (2008). Microsoft: We took out storm botnet. *Computerworld*. Retrieved August 2, 2012, from http:// www.computerworld.com/s/article/9079653/Microsoft_We_took_out_Storm_botnet

Kempa, M., Carrier, R., Wood, J., & Shearing, C. (1999). Reflections of the evolving concept of 'private policing'. *European Journal on Criminal Policy and Research*, *7*(2), 197-223. doi:10.1023/A:1008705411061

Klang, M. (2005). Free software and open source: The freedom debate and its consequences. *First Monday*, *10*(3). Retrieved from http://firstmonday.org/issues/issue10_3/klang/index.html

Klievink, B. (2011). *Unravelling interdependence — coordinating public-private service networks* [Doctoral Thesis, Delft University of Technology]. Delft: Uitgeverij BOXPress.

Klimburg, A. (2011). Mobilising cyber power. *Survival*, *53*(1), 41–60. doi:10.1080/00396338.2011.555595

Konieczny, P. (2010). Adhocratic governance in the Internet age: A case of Wikipedia. *Journal of Information Technology & Politics*, *7*(4), 263-283. doi:10.1080/19331681.2010.489408

Koot, M. R. (2012). Dutch govt expresses intent to draft new cybercrime legislation. *Matthijs R. Koot's notebook* [Web page]. Retrieved March 12, 2014, from http://blog.cyberwar.nl/2012/10/dutch-govt-expresses-intent-to-draft.html

Kosachev, K. (2007, March 6). An insult to our war dead. *The Guardian.* Retrieved from http://www.guardian.co.uk/commentisfree/2007/mar/06/comment.second-worldwar

Kostakis, V. (2010). Identifying and understanding the problems of Wikipedia's peer governance. The case of inclusionists versus deletionists. *First Monday*, *15*(3). Retrieved from http://firstmonday.org/ojs/index.php/fm/article/view/2613

Kostakis, V. (2011). Commons-based peer production and the neo-weberian state: Synergies and interdependencies. *Halduskultuur-Administrative Culture*, *12*(2), 146-161.

Koszarek, W. E. (2009). *Peer production in the US Navy: Enlisting Coase's penguin.* Master thesis. Retrieved October 28, 2013, from http://www.dtic.mil/cgi-bin/GetTRDoc?AD=ADA514470

Krahmann, E. (2003). Conceptualizing security governance. *Cooperation and Conflict*, *38*(1), 5. doi:10.1177/0010836703038001001

Krahmann, E. (2005). Security governance and networks: New theoretical perspectives in transatlantic security. *Cambridge Review of International Affairs*, *18*(1), 15-30. doi:10.1080/09557570500059514

Krahmann, E. (2007). Risk markets: The commodification of security and the risk society. [ECPR Standing Group on International Relations (SGIR), 6th Pan-European International Relations Conference, Turin, 12-15 September 2007]. Retrieved from http://www.eisa-net.org/bebruga/eisa/files/events/turin/Elke-krahmann_turin_paper_final.pdf

Krahmann, E. (2008). Security: Collective good or commodity? *European Journal of International Relations*, *14*(3), 379-404. doi:10.1177/1354066108092304

Krahmann, E. (2010). *States, citizens and the privatization of security.* Cambridge, UK, New York: Cambridge University Press.

Krebs, B. (2006, March 21). Bringing botnets out of the shadows – online volunteers monitor illegal computer networks. *Washington Post.* Retrieved January 2, 2013, from http://www.washingtonpost.com/wp-dyn/content/article/2006/03/21/AR2006032100279_pf.html

Krebs, B. (2009, March 16). Massive profits fueling rogue antivirus market. *Washington Post.* Retrieved February 10, 2013, from http://voices.washingtonpost.com/securityfix/2009/03/obscene_profits_fuel_rogue_ant.html

Krebs, B. (2011). $72M scareware ring used Conficker worm. *KrebsOnSecurity* [Web page]. Retrieved February 10, 2014, from http://krebsonsecurity.com/2011/06/72m-scareware-ring-used-conficker-worm/#more-10417

Kreiss, D., Finn, M., & Turner, F. (2011). The limits of peer production: Some reminders from Max Weber for the network society. *New Media & Society*, *13*(2), 243-259. doi:10.1177/1461444810370951

Krieger, K. (1992). *Enzyklopädie Deutscher Geschichte 14: König, Reich und Reichsreform im Spätmittelalter.* München: Oldenbourg.

Kristoff, J. (2010, November). Who + what + where — the Internet and security ops community. [Presentation given at TDC375 Spring 2010/11 onference]. Retrieved from http://condor.depaul.edu/jkristof/tdc375/2010_11-901/whowhatwhere.pdf

Kuckartz, U. (2010). *Einführung in die computergestützte Analyse qualitativer Daten* (3. ed.). Wiesbaden: VS Verlag für Sozialwissenschaften.

Kuerbis, B. (2011). *Securing critical Internet resources: Influencing Internet governance through social networks and delegation* [Doctoral dissertation]. iSchool Information Science and Technology - Dissertations, Paper 68. Retrieved from http://surface.syr.edu/it_etd/68

Kuerbis, B. (2012). European privacy authorities object to ICANN WHOIS proposals. *IGP blog* [Web page]. Retrieved March 12, 2014, from http://www.internetgovernance.org/2012/09/27/european-privacy-authorities-object-to-icann-whois-proposals/

Lake, D. A. (2009). Hobbesian hierarchy: The political economy of political organization. *Annual Review of Political Science*, *12*, 263-283. doi:10.1146/annurev.polisci.12.041707.193640

Lakhani, K. R., & Von Hippel, E. (2003). How open source software works: 'Free' user-to-user assistance. *Research Policy*, *32*(6), 923-943.

Landler, M., & Markoff, J. (2007). In Estonia, what may be the first war in cyberspace. *International Herald Tribune.* Retrieved November 4, 2010, from http://www.iht.com/articles/2007/05/28/business/cyberwar.php

Laufen, K. (2010). Cybercrime - Tatort Internet. *Das ARD Radiofeature* [Audiovisual Material]. SWR2. Retrieved February 10, 2010, from http://web.ard.de/radio/radiofeature/#awp::?page_id=117

Leder, F. (2008, December 30). Stormfucker. [leder@informatik.uni-bonn.de, email sent to Full disclosure mailing list]. Retrieved September 17, 2011, from http://seclists.org/fulldisclosure/2008/Dec/588

Leder, F., & Werner, T. (2009). *Containing Conficker. To tame a malware* (rev1 ed.) The Honeynet Project. KYE Paper: Know your enemy. Retrieved December 10, 2009, from http://www.honeynet.org/files/KYE-Conficker.pdf

Lee, E. (2003). The public's domain: The evolution of legal restraints on the government's power to control public access through secrecy or intellectual property. *Hastings Law Journal*, *55*, 91-210.

Leiner, B. M., Cerf, V. G., Clark, D. D., Kahn, R. E., Kleinrock, L., Lynch, D. C., . . . Wolff, S. (2012). A brief history of the Internet. [Web page] Internet Society. Retrieved August 13, 2013, from http://www.internetsociety.org/brief-history-internet

Leonard, A. (2002). The end of the revolution. *Salon* [Web page]. Retrieved from http://dir.salon.com/story/tech/books/2002/06/14/root/print.html

Leyden, J. (2004, May 17). Phatbot suspect released on bail. *The Register.* Retrieved February 16, 2014, from http://www.theregister.co.uk/Print/2004/05/17/phatbot_suspect_bailed/

Leyden, J. (2010, March 11). Estonian DDoS revenge worm crafter jailed—infection still spreading. *The Register.* Retrieved February 23, 2014, from http://www.theregister.co.uk/2010/03/11/allaple_ddos_vxer_jailed/

Libicki. (2009). *Cyberdeterrence and cyberwar.* RAND. Retrieved from http://www.rand.org/pubs/monographs/MG877/?ref=homepage&key=t_MG877_cover

Loader, I., & Walker, N. (2007). *Civilizing security.* Cambridge, New York: Cambridge University Press.

Loose Lips Sink Ships. (2001). *Unifying A nation—World War II posters from the New Hampshire state library* [Catalogue of the exhibition "Unifying a Nation", opened in Dec 7, 2001] [Web page]. Retrieved from http://www.nh.gov/nhsl/ww2/loose.html

Loubser, M. (2006). Peer production of security information. In *Oxford university computing laboratory programming research group student conference'06* (pp. 31-32).

Loubser, M. (2008). Governance structures in distributed problem solving networks. [Oxford Internet Institute DPSN Working Paper Series No. 16]. Retrieved May 3, 2010, from http://ssrn.com/abstract=1302945

Lowry, R. P. (1972). Toward a sociology of secrecy and security systems. *Social Problems*, *19*(4), 437-450.

Luhmann, N. (2014). *Vertrauen: Ein Mechanismus der Reduktion sozialer Komplexität* (5. ed.). Konstanz, München: UVK Verlagsgesellschaft. (Original work published 1968)

Machiavelli, N. (1910). *The prince* (Project Gutenberg ebook ed.). PF Collier. (Original work published 1505) Retrieved from http://www.gutenberg.org/files/1232/1232-h/1232-h.htm

Mandelbaum, M. (2006). *The case for Goliath: How America acts as the world's government in the twenty-first century.* PublicAffairs.

Markle Foundation - Task Force on National Security in the Information Age. (2002). *Protecting America's freedom in the information age. A report of the Markle Foundation Task Force.* New York: Markle Foundation.

Markle Foundation - Task Force on National Security in the Information Age. (2003). *Creating a trusted network for homeland security. Second report of the Markle Foundation Task Force* [ee]. New York: Markle Foundation.

Markoff, J. (2009, March 18). Computer experts unite to hunt worm. *New York Times*. Retrieved March 2, 2011, from http://www.nytimes.com/2009/03/19/technology/19worm.html

Markoff, J. (2009, January 22). Worm infects millions of computers worldwide. *New York Times.* Retrieved March 2, 2011, from http://www.nytimes.com/2009/01/23/technology/internet/23worm.html

Mathew, A., & Cheshire, C. (2010). The new cartographers: Trust and social order within the Internet infrastructure. In *TPRC 2010.* Retrieved from http://ssrn.com/abstract=1988216

Mathiason, J., Mueller, M., Klein, H., Holitscher, M., & McKnight, L. (2004). *Internet governance: The state of play*. Retrieved February 23, 2008, from http://www.internetgovernance.org/pdf/ig-sop-final.pdf

Maxwell, E. (2006). Open standards, open source, and open innovation: Harnessing the benefits of openness. *Innovations: Technology, Governance, Globalization*, *1*(3), 119-176. Retrieved April 25, 2010, from http://www.mitpressjournals.org/doi/pdfplus/10.1162/itgg.2006.1.3.119

Mayer-Schonberger, V. (2009). Can we reinvent the Internet? *Science*, *325*(5939), 396.

Mayntz, R. (2004). Organizational forms of terrorism: Hierarchy, network, or a type sui generis? [MPIfG Discussion Paper 04/4]. Köln: Max Planck Institute for the Study of Societies.

McAfee. (2008). Exploit-MS08-067 bundled in commercial malware kit. *McAfee blog central* [Web page]. Retrieved February 11, 2014, from http://blogs.mcafee.com/mcafee-labs/exploit-ms08-067-bundled-in-commercial-malware-kit

McChesney, F. S. (1986). Government prohibitions on volunteer fire fighting in nineteenth-century America: A property rights perspective. *The Journal of Legal Studies*, *15*(1), 69-92.

Meeker, M. (2012, May 30). Internet trends. [Presentation given at D10 Conference]. Retrieved December 5, 2012, from http://www.kpcb.com/insights/2012-internet-trends

Meiroșu, C. (2007a). Meeting minutes, 21st TF-CSIRT meeting, May 3-4, 2007, Prague, Czech Republic. TERENA. Retrieved February 22, 2014, from http://www.terena.org/activities/tf-csirt/meeting21/21st_tfcsirt_meeting_minutes.pdf

Meiroșu, C. (2007b). Meeting minutes, 22nd TF-CSIRT meeting, 20-21 September 2007, Porto, Portugal. TERENA. Retrieved February 22, 2014, from http://www.terena.org/activities/tf-csirt/meeting22/22nd_tfcsirt_meeting_minutes.pdf

Menn, J. (2010). *Fatal system error: The hunt for the new crime lords who are bringing down the Internet.* New York: PublicAffairs.

Mezgar, I. (2003). Role of trust in networked production systems. *Annual Reviews in Control*, *27*(2), 247-254. doi:10.1016/j.arcontrol.2003.09.007

Microsoft. (2009a). *Microsoft Security Intelligence Report, January through June 2009* (Vol. 7). Retrieved December 29, 2009, from www.microsoft.com/sir

Microsoft. (2009b). *Microsoft Security Intelligence Report, July through December 2008* (Vol. 6). Retrieved December 29, 2009, from www.microsoft.com/sir

Microsoft. (2012). *Microsoft Security Intelligence Report, July through December 2011* (Vol. 12). Retrieved from www.microsoft.com/sir

Microsoft. (2013). *Microsoft Security Intelligence Report, January to June 2013* (Vol. 15). Microsoft. Retrieved from www.microsoft.com/sir

Microsoft Corporation. (2008). Microsoft security bulletin MS08-067. *Microsoft Technet* [Version 1.0, Oct. 23]. Retrieved October 8, 2010, from http://www.microsoft.com/technet/security/bulletin/ms08-067.mspx

Microsoft SDL Team. (2008). MS08-067 and the SDL. *MSDN security development lifescycle blog* [Web page]. Retrieved February 11, 2014, from http://blogs.msdn.com/b/sdl/archive/2008/10/22/ms08-067.aspx

Miller, R. A., & Kuehl, D. T. (2009). Cyberspace and the 'first battle' in 21st-century war. *Defence horizons* (Vol. 68, Sept. 2009). National Defense University, Center for Technology and National Security Policy.

Ministerie van Veiligheid en Justitie. (2012, October 15). Aan de voorzitter van de tweede kamer der staten-generaal, wetgeving bestrijding cybercrime. Retrieved October 20, 2012, from http://www.rijksoverheid.nl/bestanden/documenten-en-publicaties/kamerstukken/2012/10/15/wetgeving-bestrijding-cybercrime/wetgeving-bestrijding-cybercrime-1.pdf

Moore, T. (2008). *Cooperative attack and defense in distributed networks.* Retrieved January 3, 2009, from http://www.cl.cam.ac.uk/techreports/UCAM-CL-TR-718.pdf

Moore, T., & Clayton, R. (2008). The impact of incentives on notice and take-down. *The Impact of Incentives on Notice and Take-down.* Retrieved December 28, 2008, from http://weis2008.econinfosec.org/papers/MooreImpact.pdf

Moore, T., Clayton, R., & Anderson, R. (2009). The economics of online crime. *Journal of Economic Perspectives*, *23*(3), 3-20. Retrieved from http://www.jstor.org/stable/27740537

Morozov, E. (2013). *To save everything, click here: The folly of technological solutionism.* New York: PublicAffairs.

Mueller, M. (2010). *Networks and states - the global politics of Internet governance.* Cambridge, Mass.: MIT Press.

Mueller, M., Ku, M., & Schmidt, A. (2009). *Transnational regimes and the response to cybercrime: The case of phishing.* Paper presented at the annual convention of the International Studies Association, New York, 14-18 Feb 2009.

Mueller, M., Schmidt, A., & Kuerbis, B. (2013). Internet security and networked governance in international relations. *International Studies Review*, (15), 86-104. doi:10.1111/misr.12024

Mueller, M. (2002). *Ruling the root: Internet governance and the taming of cyberspace.* Cambridge, Mass.: MIT Press.

Mützenich, R. (2009). Die Nutzung neuer Medien als Instrument russischer Außenpolitik in seinem „Nahen Ausland“. Website of German MP R. Mützenich [Web page]. Retrieved from http://www.rolfmuetzenich.de/texte_und_reden/veroeffentlichungen/Muetzenich_SF.pdf

Myers, S. L. (2007, January 25). Debate renewed: Did Moscow free Estonia or occupy it? *New York Times*. Retrieved from http://www.nytimes.com/2007/01/25/world/europe/25tallinn.html

Nakashima, E. (2011, May 11). Pentagon to expand cybersecurity program for defense contractors. *Washington Post*. Retrieved March 10, 2014, from http://www.washingtonpost.com/world/national-security/pentagon-to-expand-cybersecurity-program-for-defense-contractors/2012/05/11/gIQALhjbHU_story.html

National Institute of Standards and Technology. (2008). Vulnerability summary for CVE-2008-4250. *National vulnerability database, national cyber-alert system* [Last revised on June 14, 2011]

[Web page]. (Original work published October 23, 2008) Retrieved June 29, 2011, from http://web.nvd.nist.gov/view/vuln/detail?vulnId=CVE-2008-4250

NATO Sees Recent Cyber Attacks on Estonia As Security Issue. (2007, May 26). *DW-World*. Retrieved from http://www.dw-world.de/dw/article/0,,2558579,00.html

de Natris, W. (2012). Public private cooperaration: The Zeus take down example. *Personal blog* [Web page]. Retrieved January 10, 2013, from http://woutdenatris.wordpress.com/2012/05/22/public-private-cooperaration-the-zeus-take-down-example

Nazario, J. (2007). *Estonian DDoS attacks - A summary to date*. Arbor Networks. Retrieved from http://asert.arbornetworks.com/2007/05/estonian-ddos-attacks-a-summary-to-date/

Neilson, D. H. (2009). Peer production, division of labor, and project modularity. [Eastern Economic Association Meetings February 27 - March 1, 2009, New York City]. Retrieved from http://andromeda.rutgers.edu/~jmbarr/EEA2009/neilson.pdf

Newell, S., & Swan, J. (2000). Trust and inter-organizational networking. *Human Relations*, *53*(10), 1287-1328.

Nye, J. S. (1990). *Bound to lead: The changing nature of American power.* Basic Books.

Nye, J. S. (2011a). Nuclear lessons for cyber security? *Strategic Studies Quarterly*, *5*(4), 18-38.

Nye, J. S. (2011b). Power and foreign policy. *Journal of Political Power*, *4*(1), 9-24. doi:10.1080/2158379X.2011.555960

Nye, J. S., & Owens, W. A. (1996). America's information edge. *Foreign Affairs*, *2*, 20-36.

Nyswander Thomas, R. (2013). *Securing cyberspace through public-private partnership — A comparative analysis of partnership models.* Master thesis, updated August 2013. Retrieved August 25, 2013, from http://csis.org/files/publication/130819_tech_summary.pdf

Onken, E. -C. (2007). The Baltic states and Moscow's 9 may commemoration: Analysing memory politics in Europe. *Europe-Asia Studies*, *59*(1), 23-46.

Osterloh, M., & Rota, S. (2004). Trust and community in open source software production. *Analyse & Kritik*, (26), 279-301.

Osterloh, M., Rota, S., & Kuster, B. (2006). Open source software production: Climbing on the shoulders of giants. [Working Paper]. Universität Zürich - Lehrstuhl für Organisation, Technologie- und Innovationsmanagement. Retrieved October, 2008, from http://www.iou.unizh.ch/orga/downloads/publikationen/osterlohrotakuster.pdf

Ostrom, E. (2008). Institutions and the environment. *Economic Affairs*, *28*(3), 24-31. doi:10.1111/j.1468-0270.2008.00840.x

Ostrom, E., & Gardner, R. (1993). Coping with asymmetries in the commons: Self-Governing irrigation systems can work. *The Journal of Economic Perspectives*, *7*(4), 93-112. doi:10.2307/2138503

Ostrom, E., Gardner, R., & Walker, J. (1994). Institutional analysis and common-pool resources. In *Rules, games, and common-pool resources* (pp. 23-50). University of Michigan Press.

O'Toole, M. (2009). Crime, ink: There is no hunting like the hunting of man, and those who have hunted men long enough and liked it, never care for anything else thereafter. *Crime, ink (personal blog)* [Web page]. Retrieved March 12, 2014, from http://irishcrimereporter.blogspot.nl/2009/02/there-is-no-hunting-like-hunting-of-man.html

Ottis, R. (2009). Evgeny Morozov on cyber myths. *Conflicts in cyberspace* [Personal blog]. Retrieved December 12, 2011, from http://conflictsincyberspace.blogspot.com/2009/06/evgeny-morozov-on-cyber-myths.html

Ottis, R. (2010). From pitchforks to laptops: Volunteers in cyber conflicts. In C. Czosseck & K. Podins (Eds.), *Conference on cyber conflict proceedings 2010.* Tallinn, Estonia: CCD COE Publications. Retrieved from https://docs.google.com/fileview?id=0B7yq33Gize8yNDE1Zjk4ZDEtNGRhMy00YjNhLTljMjktYmEyNTc3ODc2ZDVi&hl=en

Paglen, T. (2013). *Seeing the secret state — six landscapes* [Talk given at 30C3, Hamburg, December] [Online video]. Retrieved January 2, 2014, from media.ccc.de

Parker, T. (2011, December 4). Debunking the Conficker-Iranian nuclear program connection. *Dark Reading.* Retrieved February 13, 2014, from http://www.darkreading.com/views/debunking-the-conficker-iranian-nuclear/232200687?printer_friendly=this-page

Parker, R. (2007). Networked governance or just networks? Local governance of the knowledge economy in Limerick (Ireland) and Karlskrona (Sweden). *Political Studies*, *55*(1), 113-132. doi:10.1111/j.1467-9248.2007.00624.x

Pääbo, H. (2008). War of memories: Explaining memorials war in Estonia. *Baltic Security & Defence Review*, *10*, 5-26.

Piscitello. (2010). *Conficker summary and review.* ICANN. Retrieved October 4, 2010, from https://www.icann.org/en/security/conficker-summary-review-07may10-en.pdf

Popper, K. R. (1962). *Conjectures and refutations: The growth of scientific knowledge.* Routledge & Kegan Paul.

Porras, P., Saidi, H., & Yegneswaran, V. (2009a). *An analysis of Conficker's logic and rendezvous points.* SRI International. (Original work published February 4, 2009) Retrieved October 10, 2010, from http://mtc.sri.com/Conficker

Porras, P., Saidi, H., & Yegneswaran, V. (2009b). *Conficker C analysis.* SRI International. (Original work published March 8, 2009) Retrieved August 2, 2011, from http://mtc.sri.com/Conficker/addendumC/index.html

Porras, P., Saidi, H., & Yegneswaran, V. (2009a). *Conficker C P2P protocol and implementation.* SRI International. Retrieved August 2, 2011, from http://mtc.sri.com/Conficker/P2P/

Porras, P., Saidi, H., & Yegneswaran, V. (2009b). A foray into Confickers logic and rendezvous points. [LEET '09 — USENIX Workshop on Large-Scale Exploits and Emergent Threats, Boston, MA, April 21, 2009]. Retrieved November 24, 2011, from http://static.usenix.org/event/leet09/tech/full_papers/porras/porras_html/

Poulsen, K. (2007). 'Cyberwar' and Estonia's panic attack. *Wired, threat level.* Retrieved November 10, 2010, from http://www.wired.com/threatlevel/2007/08/cyber-war-and-e/

Powell, W. W. (1990). Neither market nor hierarchy: Network forms of organization. *Research in Organizational Behavior*, *12*, 295-336.

Project Grey Goose. (2008). *Russia/Georgia cyber war – findings and analysis.* GreyLogic. Phase I report. Retrieved from http://fserror.com/pdf/GreyGoose1.pdf

Project Grey Goose. (2009). *The evolving state of cyber warfare.* GreyLogic. Phase II report. Retrieved January 9, 2011, from http://fserror.com/pdf/GreyGoose2.pdf

Randel, T. (2008). CyberWar in Estonia 2007 — history, analysis. [Presentation given at IMPACT 2008, Malaysia]. Retrieved from http://www.impact-alliance.org/downloads/TarmoRandel_CyberwarInEstonia.pdf

Rantanen, M. (2007). Virtual harassment, but for real. *Helsingin Sanomat International Edition.* Retrieved from http://www.hs.fi/english/article/Virtual+harassment+but+for+real+/1135227099868

Raustiala, K. (2002). The architecture of international cooperation: Transgovernmental networks and the future of international law. *Virginia Journal of International Law*, *43.* Retrieved July 4, 2012, from http://ssrn.com/abstract_id=333381

Raymond, E. S. (1998a). The cathedral and the bazaar. *First Monday*, *3*(3). doi:10.5210/fm.v3i2.578

Raymond, E. S. (1998b). Homesteading the noosphere. *First Monday*, *3*(10). doi:10.5210/fm.v3i10.621

Rendon Group. (2011). *Conficker Working Group: Lessons learned* [Report created in June 2010, commissioned by the Department of Homeland Security]. Retrieved January, 2011, from http://www.confickerworkinggroup.org/wiki/uploads/Conficker_Working_Group_Lessons_Learned_17_June_2010_final.pdf

Richter, E. (1992). *Theorie und Gesellschaft: Der Zerfall der Welteinheit. Vernunft und Globalisierung in der Moderne* [Habilitationsschrift]. Frankfurt/Main (u.a.): Campus.

Riden, J. (2008). How fast-flux service networks work. *The Honeynet Project* [Web page]. Retrieved February 11, 2012, from http://www.honeynet.org/node/132

Riehle, D. (2007). The economic motivation of open source software: Stakeholder perspectives. *Computer*, *40*(4), 25.

Riehle, D. (2010). The economic case for open source foundations. *Computer*, *43*(1), 93-97.

Robb, D. (2010, October 7). IBM System z buyer's guide. *ServerWatch.* Retrieved March 6, 2014, from http://www.serverwatch.com/print.php/3907336

Roberts, A. S. (2006). *Blacked out : Government secrecy in the information age* (illustrated, annotated, reprint ed., p. 322). Cambridge, New York: Cambridge University Press.

Rocco, S., Claudio, M., & Ramirez-Marquez, J. E. (2011). Vulnerability metrics and analysis for communities in complex networks. *Reliability Engineering & System Safety*, *96*, 1360-1366. doi:10.1016/j.ress.2011.03.001

Russel. (1950). What desires are politically important? [Nobel Lecture, December 11, 1950] [Web page]. Nobel Media AB. Retrieved March 11, 2014, from http://www.nobelprize.org/nobel_prizes/literature/laureates/1950/russell-lecture.html

Sales, N. A. (2007). Secrecy and national-security investigations. *Alabama Law Review*, *58*, 811.

Samuelson, L. (2006). IBM's pragmatic embrace of open source. *Communications of the ACM*, *49*, 10-21.

Schilcher, C., Poth, A. -K., Sauer, S., Stiefel, K. -P., & Will-Zocholl, M. (2011). Trust in international teams: Cultural, spatial, and organizational issues. *International Journal of Business Research*, *11*(4), 29-38.

Schiller, C. A., Binkley, J., Harley, D., Evron, G., Bradley, T., Willems, C., & Cross, M. (2007). *Botnets: The killer web app.* Syngress. Retrieved from http://sirpabs.ilahas.com/ebooks/Computer%20&%20Technology/Botnets%20-%20The%20Killer%20Web%20App%20(2007).pdf

Schmidt, A. (2004). *Hegemonie durch Technologie — Strategien zur Instrumentalisierung von Informationstechnologien in globaler Politik* [Master's thesis, Universität Kassel].

Schmidt, A. (2009). *Conceptualizing Internet security governance.* Paper prepared for Giganet conference, Sharm El-Sheik, Egypt, November 2009. Retrieved from http://netdefences.com/wp-content/uploads/Schmidt-2009-Conceptualizing-Internet-Security-Governance.pdf

Schmidt, A. (2012). At the boundaries of peer production: The organization of Internet security production in the cases of Estonia 2007 and Conficker. *Telecommunications Policy*, *36*(6), 451-461. doi:10.1016/j.telpol.2012.02.001

Schmidt, A. (2013a). The Estonian cyberattacks. In J. Healey (Ed.), *The fierce domain – conflicts in cyberspace 1986-2012.* Washington, D.C.: Atlantic Council.

Schmidt, A. (2013b). Open security. Contributions of networked approaches to the challenge of democratic Internet security governance. In R. Radu, J. Chenou, & R. Weber (Eds.), *The evolution of global Internet governance — principles and policies in the making* (pp. 169-187). Zürich: Schulthess.

Schmidt, A. (2013c). The unfolding of the information umbrella. [Personal blog]. Retrieved February 24, 2014, from https://netdefences.com/2013/06/the-unfolding-of-the-information-umbrella/

Schmidt, A. (2014). Hierarchies in networks—emerging hybrids of networks and hierarchies for producing Internet security. In J. Kremer & B. Müller (Eds.), *Cyber space and international relations—theory, prospects and challenges* (pp. 181-202). Berlin, Heidelberg: Springer. doi:10.1007/978-3-642-37481-4_11

Schmitt, C. (1985). *Political theology: Four chapters on the concept of sovereignty* [Politische Theologie. Vier Kapitel zur Lehre von der Souveränität] [Translated by George Schwab]. Cambridge, MA: MIT Press. (Original work published 1922)

Schuppert, G. F. (2008). Von Ko-Produktion von Staatlichkeit zur Co-Performance of Governance: Eine Skizze zu kooperativen Governance-Strukturen von den Condottieri der Renaissance bis zu Public Private Partnerships. *SFB-Governance Working Paper Series*, *12*.

Schweik, C. M., & English, R. (2007). Tragedy of the FOSS commons? Investigating the institutional designs of free/libre and open source software projects. *First Monday*, *12*(2). Retrieved November 8, 2008, from http://firstmonday.org/issues/issue12_2/schweik/index.html

Schweik, C. M., English, R., & Haire, S. (2008a). Factors leading to success or abandonment of open source commons: An empirical analysis of sourceforge.net projects. *South African Computer Journal*, (43), 58-65. Retrieved January 4, 2010, from http://works.bepress.com/charles_schweik/15/

Schweik, C. M., English, R., & Haire, S. (2008b). Open source software collaboration: Foundational concepts and an empirical analysis. *National Center for Digital Government Working Paper Series*, *8*(2), 1. Retrieved from http://scholarworks.umass.edu/ncdg/28/

Segura, V., & Lahuerta, J. (2010). Modeling the economic incentives of DDoS attacks: Femtocell case study. In T. Moore, D. J. Pym, & C. Ioannidis (Eds.), *Economics of information security and privacy* (pp. 107-119). New York et al. : Springer. doi:10.1007/978-1-4419-6967-5_7

Shakespeare, W. (2000). *Hamlet, Prince of Denmark* [The Oxford Shakespeare: the complete works of William Shakespeare, ed. W.J. Craig]. London: Bartleby.com. (Original work published 1601) Retrieved from http://www.bartleby.com/70/index42.html

Shapiro, F. R. (1987). Etymology of the computer bug: History and folklore. *American Speech*, *62*(4), 376-378. Retrieved from http://www.jstor.org/stable/455415 .

Shearing, C., & Wood, J. (2003). Nodal governance, democracy, and the new 'denizens'. *Journal of Law and Society*, *30*(3), 400-419. doi:10.1111/1467-6478.00263

Shirky, C. (2005). Epilogue: Open source outside the domain of software. *Perspectives on Free and Open Source Software*, 483-488.

Shirky, C. (2008). *Here comes everybody: The power of organizing without organizations.* Penguin Press.

Shnayerson, M. (2004, January). The code warrior. *Vanity Fair.* Retrieved June 1, 2011, from http://www.vanityfair.com/culture/features/2004/01/virus-hunters-200401

Sietmann, R. (2005, December 29). 22C3: Pro und kontra e-Voting. *heise online.* Retrieved January 15, 2011, from http://www.heise.de/newsticker/meldung/22C3-Pro-und-Kontra-e-Voting-161678.html

Singel, R. (2008, February 25). Pakistan's accidental YouTube re-routing exposes trust flaw in net. *Wired.* Retrieved March 11, 2014, from http://www.wired.com/threatlevel/2008/02/pakistans-accid/

Singel, R. (2010, March 1). Cyberwar hype intended to destroy the open Internet. *Wired - Threat Level.* Retrieved February 13, 2014, from http://www.wired.com/threatlevel/2010/03/cyber-war-hype/

Singh, J. P. (2002). Introduction: Information technologies and the changing scope of global power and governance. In J. N. Rosenau & J. P. Singh (Eds.), *Information technologies and global politics. The changing scope of power and governance.* Albany (NY): State University of New York Press.

Skolnikoff, E. B. (1993). *The elusive transformation. Science, technology, and the evolution of international politics.* Princeton: Princeton University Press.

Slaughter, A. (2011). *Lego world* [Presentation at PopTech conference 2011] [Video]. Retrieved July, 2013, from http://poptech.org/popcasts/annemarie_slaughter_lego_world

Slaughter, A. M. (2004). *A new world order.* Princeton, NJ: Princeton University Press.

Slaughter, A. M. (2009). America's edge: Power in the networked century. *Foreign Affairs*, *88*(1), 94-113.

Socor, V. (2007). Moscow stung by Estonian ban on totalitarianism's symbols. *Eurasia Daily Monitor, the Jamestown Foundation.* Retrieved from http://jamestown.org/single/?no_cache=1&tx_ttnews[tt_news]=32427

Solomon, A., & Evron, G. (2006). The world of botnets. *Virus Bulletin*, 10-12. Retrieved from http://www.gadievron.com/publications/SolomonEvronSept06.pdf

Sotirov, A. (2008). Decompiling the vulnerable function for MS08-067. *Security research (personal blog)* [Web page]. Retrieved February 11, 2014, from http://www.phreedom.org/blog/2008/decompiling-ms08-067/

Stanković, S., & Simić, D. (2009). Defense strategies against modern botnets. *International Journal of Computer Science and Information Security*, *2*(1). Retrieved December 12, 2009, from http://arxiv.org/abs/0906.3768

Steigerwald, D., Vigna, G., Kruegel, C., Kemmerer, R., Abman, R., & Stone-Gross, B. (2011). The underground economy of fake antivirus software. *Departmental working papers*. Department of Economics, UC Santa Barbara. Retrieved March 3, 2012, from http://escholarship.org/uc/item/7p07k0zr

Sternstein, A. (2012, May 11). Pentagon opens classified cyber program to all defense contractors, ISPs. *Forbes.* Retrieved from http://www.forbes.com/sites/davidthier/2012/04/09/the-fight-for-the-internet-continues-could-cispa-be-the-next-sopa/

Stockton, P. N., & Golabek-Goldman, M. G. (2013). Curbing the market for cyber weapons. *Yale Law & Policy Review, Forthcoming.* Retrieved from http://ssrn.com/abstract=2364658

Stone, B. (2010). Google's Andy Rubin on everything Android. *New York Times* [Web page]. Retrieved September 26, 2013, from http://bits.blogs.nytimes.com/2010/04/27/googles-andy-rubin-on-everything-android

Stone-Gross, B., Holz, T., Stringhini, G., & Vigna, G. (2011). The underground economy of spam: A botmaster's perspective of coordinating large-scale spam campaigns. In *USENIX workshop on large-scale exploits and emergent threats (LEET).*

Suganami, H. (2007). Understanding sovereignty through Kelsen/Schmitt. *Review of International Studies*, *33*(03), 511-530. doi:10.1017/S0260210507007632

Sun Tzu. (2009). *The art of war* [translated by Lionel Giles]. Pax Librorum Publishing House. (Original work published 1910)

Swire, P. (2004). *A model for when disclosure helps security: What is different about computer and network security?*. Moritz College of Law, Ohio State University. Center for Law, Policy and Social Science – Working Paper Series, No. 12. Retrieved May 1, 2010, from http://ssrn.com/abstract=531782

Swire, P. P. (2005). Theory of disclosure for security and competitive reasons: Open source, proprietary software, and government systems. *Houston Law Review*, *42*, 1333-1380.

Symantec. (2009). *The Downadup codex. A comprehensive guide to the threat's mechanics* (2.0 ed.). Retrieved October 10, 2010, from http://www.symantec.com/content/en/us/enterprise/media/security_response/whitepapers/the_downadup_codex_ed2.pdf

Sørensen, E., & Torfing, J. (2007). *Theories of democratic network governance.* Palgrave Macmillian.

Tabatabaie, S., van Eeten, M., & Asghari, H. (2012). Transgovernmental networks in cybersecurity: A quantitative analysis of the London Action Plan against spam. [2012 Annual Convention of the International Studies Association].

Thompson, P. B. (2000). *Privacy, secrecy and security* (pp. 0-7). Center for Education and Research in Information Assurance and Security, Purdue University, West Lafayette. CERIAS Tech Report 2000-18.

Thomson, I. (2007, May 31). Russia 'hired botnets' for Estonia cyber-war - Russian authorities accused of collusion with botnet owners. *Computing.co.uk.* Retrieved from http://www.computing.co.uk/vnunet/news/2191082/claims-russia-hired-botnets

Tikk, E. (2010). Global cybersecurity — thinking about the niche for NATO. *SAIS Review*, *30*(2), 105-119.

Tikk, E., Kaska, K., & Vihul, L. (2010). *International cyber incidents — legal considerations.* Tallinn, Estonia: Cooperative Cyber Defence of Excellence (CCD COE).

Tilly, C. (1985). War making and state making as organized crime. In P. Evans, D. Rueschemeyer, & T. Skocpol (Eds.), *Bringing the state back in* (pp. 169-191). Cambridge: Cambridge University Press.

Traynor, I. (2007). Russia accused on unleashing cyberwar to disable Estonia. *The Guardian.* Retrieved from http://www.guardian.co.uk/world/2007/may/17/topstories3.russia

Trnka, J., & Johansson. (2009). Collaborative command and control practice: Adaptation, self-regulation and supporting behavior. *International Journal of Information Systems for Crisis Response and Management*, *1*(2), 47-67.

Tsiakis, T., & Sthephanides, G. (2005). The concept of security and trust in electronic payments. *Computers & Security*, *24*(1), 10-15.

U.S. Joint Chiefs of Staff. (1997). *Joint doctrine for operations security* [Joint Pub 3-54].

U.S. Marine Corps. (2001). *Information operations* (Coordinating Draft 2-27-01 ed.) [Marine Corps Warfighting Publication (MCWP) 3-36].

UCLA Center for Digital Humanities. (2008). A digital humanities manifesto. [Web page] Retrieved March 11, 2014, from http://manifesto.humanities.ucla.edu/2008/12/15/digital-humanities-manifesto/

UK Launches New Cyber Security Strategy. (2011). *Computer Fraud & Security*, *2011*(12), 3. doi:10.1016/S1361-3723(11)70119-1

United States General Accounting Office. (2013). Morris worm: Virus highlights need for improved Internet management. In J. Healey (Ed.), *The fierce domain – conflicts in cyberspace 1986-2012* (pp. 107-119). Washington, D.C.: Atlantic Council. (Original work published June, 1989)

US Embassy Cables: Germany Behind NATO Proposal for Baltic States. (2010, December 6). *The Guardian*. Retrieved from http://www.guardian.co.uk/world/us-embassy-cables-documents/240187

US Embassy Tallinn. (2007a). Estonia's cyber attacks: Lessons learned. *Secret US embassy cables* [Cable 07TALLINN375 sent to US Department of States on June 6]. Wikileaks. Retrieved May 31, 2014, from https://wikileaks.org/cable/2007/06/07TALLINN375.html

US Embassy Tallinn. (2007b). Estonia's cyber attacks: World's first virtual attack against nation state. *Secret US embassy cables* [Cable 07TALLINN366 to Department of State on Jun 4]. Wikileaks. Retrieved January 14, 2011, from https://wikileaks.org/cable/2007/06/07TALLINN366.html

van Eeten, M. (2008). Internet security: Incentives, externalities and the opportunity costs of disaster resistance. [Surviving Future Disasters conference, Louisiana State University, Baton Rouge, April 6-8, 2008]. Retrieved February 2, 2010, from http://www.bus.lsu.edu/centers/sdmi/files/SFDPapers/MvEeten-SDMIConferenceBaton%20Rouge.pdf

van Eeten, M., & Bauer, J. (2008a). *ITU study on the financial aspects of network security: Malware and spam* [ICT Applications and Cybersecurity Division, Policies and Strategies Department, ITU Telecommunication Development Sector]. Retrieved December 10, 2009, from http://www.itu.int/ITU-D/cyb/cybersecurity/docs/itu-study-financial-aspects-of-malware-and-spam.pdf van Eeten, M., & Bauer, J. M. (2008b). *Economics of malware: Security decisions, incentives and externalities*. OECD publishing. OECD Science, Technology and Industry Working Papers 2008/1. doi:10.1787/241440230621

van Eeten, M., & Bauer, J. M. (2009). Emerging threats to Internet security: Incentives, externalities and policy implications. *Journal of Contingencies and Crisis Management*, *17*(4), 221-232. Retrieved from http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1508844

van Eeten, M., Asghari, H., Bauer, J. M., & Tabatabaie, S. (2011). *Internet service providers and botnet mitigation – A fact-finding study on the dutch market*.

van Eeten, M., Bauer, J., Asghari, H., Tabatabaie, S., & Rand, D. (2010). *The role of Internet service providers in botnet mitigation an empirical analysis based on spam data*.

van Eeten, M., Bauer, J., Asgharia, H., Tabatabaie, S., & Rand, D. (2012). The role of internet service providers in botnet mitigation an empirical analysis based on spam data. *Available at SSRN 1989198*.

van Wendel de Joode, D. J. R. V., de Bruijn, J. A., & van Eeten, M. J. G. V. (2003). *Protecting the virtual commons: Self-Organizing open source and free software communities and innovative intellectual property regimes* [Information Technology & Law Series, 3]. The Hague: T.M.C. Asser Press.

van Wendel de Joode, R. (2005). *Understanding open source communities—an organisational perspective* [Doctoral Thesis, Delft University of Technology, Faculty of Technology, Policy and Management]. Delft.

Venables, J. (1991). End of the peer production? *Physics World*, 16.

Venemaa Keeldub Endiselt Koostööst Küberrünnakute Uurimisel. (2008, December 13). *ERR*. Retrieved from http://uudised.err.ee/index.php?06147571

Victory Day (9 May). (2012). *Wikipedia: The free encyclopediaia*. Retrieved October 26, 2012, from http://en.wikipedia.org/w/index.php?title=Victory_Day_(9_May)&oldid=502383694

Viégas, F. B., Wattenberg, M., & McKeon, M. M. (2007). The hidden order of Wikipedia. *Lecture Notes in Computer Science*, *4564*, 445.

Vignoli, I. (2013). Getting close to LibreOffice 4.1. *The Document Foundation blog* [Web page]. Retrieved March 12, 2014, from http://blog.documentfoundation.org/2013/07/22/getting-close-to-libreoffice-4-1/

von Hippel, E. (2002). Open source software projects as user innovation networks. [Open Source Software: Economics, Law and Policy, Toulouse, France, June 20-21]. Retrieved from http://www.idei.fr/display.php?a=2493f

Von Hippel, E., & Von Krogh, G. (2003). Open source software and the" private-collective" innovation model: Issues for organization science. *Organization Science*, 209-223.

Vymětal, P. (2007). Networked governance. In J. Němec (Ed.), *Global and regional governance – Europe and beyond (6th CEEISA convention: Selected papers)* (pp. 115-125). Prague: University of Economics, Prague, Faculty of International Relations.

Wagner, R. P. (2003). Information wants to be free: Intellectual property and the mythologies of control. *Columbia Law Review*, 995-1034.

Wang, J. (2007). *The role of social networks in the success of open source systems: A theoretical framework and an empirical investigation.* Doctoral dissertation, Kent State University Graduate School of Management. Retrieved from http://rave.ohiolink.edu/etdc/view?acc_num=kent1180640710

Ward, M. D., Stovel, K., & Sacks, A. (2011). Network analysis and political science. *Annual Review of Political Science*, *14*(1), 245-264. doi:10.1146/annurev.polisci.12.040907.115949

Was It Worth It? (2008, May 1). *Baltic Times*. Retrieved October 23, 2012, from http://www.baltictimes.com/news/articles/20360/

Waugh, W. L., Jr., J., & Sylves, R. T. (2002). Organizing the war on terrorism. *Public Administration Review*, *62*, 145-153.

Weber, M. (1972). *Wirtschaft und Gesellschaft: Grundriß der verstehenden Soziologie* (Studienausgabe, 5., rev. ed.). Tübingen: Mohr Siebeck. (Original work published 1921)

Weber, S. (2000). The political economy of open source software. *California, June.* Retrieved January 12, 2010, from http://www.escholarship.org/uc/item/3hq916dc.pdf

Weber, S. (2004). *The success of open source.* Cambridge, MA: Harvard University Press.

West, J., & Gallagher, S. (2006). Patterns of open innovation in open source software. In H. Chesbrough, W. Vanhaverbeke, & J. West (Eds.), *Open innovation: Researching a new paradigm* (pp. 82-106). Oxford: Oxford University Press.

Wicherski, G., Werner, T., Leder, F., & Schlösser, M. (2008, December 29). Stormfucker: Owning the storm botnet. [Presentation given at the 25th Chaos Communication Congress, Berlin]. Retrieved September 16, 2011, from http://media.ccc.de/browse/congress/2008/25c3-3000-en-stormfucker_owning_the_storm_botnet.html

Williams, C. (2009, January 20). Conficker seizes city's hospital network. *The Register.* Retrieved February 13, 2014, from http://www.theregister.co.uk/2009/01/20/sheffield_conficker/

Willsher, K. (2009, February 7). French fighter planes grounded by computer virus. *The Telegraph.* Retrieved February 13, 2014, from http://www.telegraph.co.uk/news/worldnews/europe/france/4547649/French-fighter-planes-grounded-by-computer-virus.html

Wilson, C. (2007). *Botnets, cybercrime, and cyberterrorism: Vulnerabilities and policy issues for congress.* Congressional Research Service. CRS Report for Congress.

Woods, D. (2009, August 29). The myth of crowdsourcing. *Forbes.* Retrieved November 20, 2009, from http://www.forbes.com/2009/09/28/crowdsourcing-enterprise-innovation-technology-cio-network-jargonspy.html

World War II Casualties. (2012). *Wikipedia: The free encyclopedia.* Retrieved October 26, 2012, from http://en.wikipedia.org/wiki/World_War_II_casualties

Zedner, L. (2003). Too much security? *International Journal of the Sociology of Law*, *31*(3), 155-184.

Summary

1. Introduction

Internet security and the security of information systems are no longer narrow technological subjects, but are at the top of policy agendas worldwide. The rise of the Internet has brought new security risks. At the same time, it allows for new forms of collaboration among globally distributed teams, jointly producing intangible goods that require little to no initial capital investment and production costs.

Existing research and theories on open source and peer production have focused on describing the enabling factors of this mode of decentralised, distributed, and non-hierarchical form of collaboration. Yet, the limits of this form of production have yet to be explored, or inquiry made into possible organisational hybrids.

This study aims to contribute to the literature on open source, peer, and social production by analysing its limitations and feasibility in the domain of Internet security governance, and responses to large-scale Internet security incidents. In a way, this study places peer production and internet security in a room together for the first time, and sees how well they get along.

The first research question guiding this study is: Can the handling of a particular Internet security incident be classified as peer production? A second set of questions addresses a) the role and importance of secrecy and the antagonists in incident responses and b) factors supporting either the absence or presence of elements of peer production in response activities.

2. Theoretical Foundations

To address these questions, this study employs theories, models, and concepts like open source and peer production, Internet governance, security governance, trust and secrecy.

The idea of peer production describes collaborative arrangements for the development of free/open source software and Wikipedia. According to Benkler, this "new modality of organizing production" is facilitated by cheap-Internet based communication, distributed ownership of cheap ICT systems, and thereby reduced costs of distributed collaboration.

The study proposes a taxonomy to clearly label varieties of social production. *Social production* is the umbrella term, describing new non-market and non-hierarchy forms of production. *Peer production* requires that social production is based on a egalitarian distribution of authority among the contributors. *Commons-based peer production* labels that subset of social production, in which no central controlling governance unit or appropriation of produced goods exist.

The defining characteristics of these variants of social production are distributiveness, openness and socialness. *Distributiveness* refers to the network topology of the contributing agents, the absence of central or decentralised hubs, the prevalence of peer governance and ad-hoc hierarchies. *Openness* describes the accessibility of the production platform, its transparency, a shared sentiment, and the accessibility and modifiability of produced goods. *Socialness* refers to the social transaction framework, i.e., contributors participate not based on monetary incentives or hierarchical pressure, but on intrinsic incentives.

Untangling the relations between secrecy, security, and social production, this study provides an analytical model to explain the presence and absence of elements of social production in incident response. Some degree of secrecy is compatible with peer production, but it alters the underlying economics of the peer production model and decreases the viability and ideational ambition of the production model.

3. Research Design

This research employs the study of cases of Internet security incidents. This allows for a detailed description of actors involved in the responses, the organisation of their activities, their access to relevant information, and their interactions and collaboration.

As criteria for the selection of the cases, the incidents need to be significant and limited in time and scope. In addition, data on the incidents and the response measures are available and accessible. Eventually, the Estonian cyberattacks and the Conficker botnet were chosen.

Identifying the application of peer production in incident response requires three steps. The first task is to identify the goods and services produced in response activities, requiring a descriptive narrative thereof. Second, the activities within this response are categorized, using the specified criteria of peer production — distributiveness, openness, and socialness. The third task is to decide whether the response can be classified as peer production or not.

In order to analyse why elements of peer production were or were not applied in the responses, a model of the feasibility of peer production was developed. It describes factors that have been identified as likely prerequisites for the feasibility of peer production for the creation of a particular good.

The main source of empirical data is a series of qualitative expert interviews, supplemented by desk-research and a bit of direct observation. Interviews have been transcribed and coded.

4. Endangering the Internet

The ground for subsequent analyses was prepared with a historiographic depiction of the two incidents.

For three weeks from April 27 until May 18, 2007, components of the Estonian Internet infrastructure were overwhelmed by Distributed Denial of Service (DDoS) attacks, website defacements, DNS server attacks, mass e- mail and comment spam. The attacks constituted a watershed in the history of Internet security because of two aspects. Firstly, the attacks made it plausible to a wider public that cyberattacks could be used as a tool in international or bilateral conflicts. Secondly, the Estonian cyberattacks are a rare case where a “national security situation” was straightened out by a community of technical experts.

In late 2008, a malware exploited a critical vulnerability within the Microsoft Windows operating system, and installed itself rapidly and silently on millions of PCs. Infected computers became part of a huge botnet, created by the use of a stunning range of innovative attack techniques. Despite its unusual size, the botnet has only been used in minor cybercrime cases, making its underlying purpose mysterious to this day. The Conficker case raised awareness of the problem of botnets among a wider global audience. It also featured an impressive global collaboration among technical experts and shed light on the Internet’s commonly hidden security institutions.

5. Producing Internet Security

How is the Internet secured and defended in times of a crisis, and which products and services are provided by the responding actors to eventually re-establish the status quo ante?

The Estonian attack was mitigated by the Estonian community of technical experts, supported by their international peers. CERT-EE evolved as the central hub for information exchange and coordinated some of the defensive measures. Security production consisted of a successful mitigation of different attack techniques, DDoS mitigation first among them. Situational awareness was established by monitoring technical and social networks. Collaboration among domestic actors was eased by familiarity among the country's security professionals. With global security-communities, ad-hoc collaboration had to be established in a rather improvised manner.

Initially, the response to the Conficker botnet consisted only of a small ad-hoc group of security professionals. The technical nature of the attacks, of involved Internet subsystems, and the response strategy chosen, required a large response community. Its creation was enabled by half a decade of prior conferencing and networking. The primary response activities included malware and botnet analyses, defensive DNS, and sinkholing. Reverse engineering of the malware's binary code was indispensable to understand the botnet and design its mitigation. Through the implementation of Defensive DNS, bots were denied from contacting the botnet's command servers. Sinkholes, databases containing information about bots' traffic, were created to learn about the bots' activities and distribution.

Both responses depended on contributions from Internet security communities. Their values, norms, and internal practices influence how Internet security is produced after such large-scale incidents.

6. Social Dimensions of Internet Security

The empirical analysis identifies the role of peer production in Internet security incident response by puzzling out whether the responses to the Estonian 2007 and the Conficker incidents have been distributed, open, and social. In both cases, incident response did not conform to all aspects of the peer production model.

In terms of distributiveness, both responses blend elements of decentrality and distributiveness. Both CERT EE and the Conficker Cabal had a central role in coordinating the responses. Activities like situational monitoring, DDoS mitigation, and malware analysis were mostly distributed, while defensive DNS or traffic filtering were conducted in a decentralised manner. Individual authority was not evenly distributed within response communities. Equipotentiality among actors largely existed in both responses, though some internal hierarchies emerged. Response networks were able to enforce norms among their members.

The responses in both cases have not been open in the way that open source production is. At best, the activities happened in an environment that could be described as gated openness or barcamps within walls. Access to security communities is usually restricted, requires vetting and vouching of potential members and is based upon interpersonal trust. The response networks, however, were also comprised of unvetted actors. Within the boundaries of the security communities, many informational resources were shared. While some of them had flirted with openness ideology in the past, the guiding principle by now is responsible disclosure and responsibility towards victims.

The responses came closest to peer production in the dimension of socialness. The motivations of contributors resembled those common in open source projects, including the aim to foster idealistic values or to follow personal interests. Bringing down the 'bad guys' and 'playing with secrets' may be motivations unique for security communities, but still fit into the open source motivation framework. The same holds true for the shared interest in gaining from indirect appropriation. Furthermore, contributors have substantial commitments towards their communities, and their loyalty to communities can trump that of their employers.

7. Limits of Openness

The open source access-for-all policy and the swift trust model is replaced by a system of limited access, vetting of potential contributors and sharing on a need-to-know basis. This outcome — no pure peer production, but significant application of elements of peer production — raises the question of why certain elements of peer production can be found in the response endeavour, while others have not been applied.

Analysing the hindrances of openness based on a model of factors supporting open or secretive approaches, the study identifies the communities' antagonist, the "bad guys", as a main driver towards secrecy and the communities' preference for a walled organisational approach. The flavour of social production used in the response endeavours resembles an institutional design that tries to incorporate some major advantages of the ideal-type peer or open source production model, while at the same time factoring in the need for secrecy.

The application of deep-trust as a prerequisite for community membership can therefore be interpreted as a hedge against the risk of defections. The observed hybrid flavour of social production reduces the risks of intrusion by malevolent bad guys, who seek to nullify the communities' defence efforts. Community-based vetting and a certain degree of permeability towards new contributors keep the costs of secrecy relatively low.

While secrecy thwarts one source of peer production effectiveness — the unplanned, unrestricted use of resources by high numbers of agents with diverse talents and skills — security communities can still leverage relatively low-cost permeability to new contributors to take advantage of external information gains.

8. Conclusion

The production of Internet security looked different in the cases analysed than in a usual circumstance in which public security is affected. With transnational cooperation among police, law enforcement, or military forces lacking or inappropriate, distributed bottom-up collaboration in ad-hoc teams or permanent security communities has played a decisive role. Unlike in open source and rigid peer production, access to production platforms, input resources, and to intermediary informational goods is restricted, and no culture of unconditional openness and transparency exists.

Naturally, this study has a number of limitations. It only offers a glimpse into the relationship between peer production and Internet security. The observations allow no clear conclusions about optimal organisational designs of response endeavours. This holds even more as the response organisations had an element of historic contingency. In addition, employee-employer relationships are not based on intra-organisational data.

A number of research gaps have been observed in the course of this study. Internet security communities deserve further analyses from different theoretical angles and levels of analysis — be it Ostrom's common-pool communities, epistemic communities, or International Relations theories. More encompassing theories of trust-based collaboration, social production, distributed, networked, and egalitarian communities would be valuable. Finally, deeper theoretical and design studies on open security systems are recommended.

The study concludes with discussions on the state of Internet security governance and ideas on how to 'open' it. More recent trends in Internet security governance have nurtured the impression of a relative decline in Internet security communities. On the other hand, centralising effects and the hierarchification of the community could be avoided by a range of measures.

Samenvatting (Dutch summary)

Geheimhouding versus openheid — internetveiligheid en de grenzen van open source- en peer-productie

1. Introductie

Internetveiligheid en de veiligheid van informatiesystemen zijn niet langer onderwerpen die alleen technologisch interessant zijn, maar wereldwijd bovenaan beleidsagenda's staan. De opkomst van het internet heeft nieuwe veiligheidsrisico's met zich meegebracht. Tegelijkertijd maakt het internet nieuwe vormen van samenwerking mogelijk tussen teams die verspreid over de wereld gezamenlijk werken aan goederen en diensten die geen of weinig startkapitaal en productiekosten vereisen.

Bestaand onderzoek en theorieën over open source- en peer-productie hebben zich geconcentreerd op het beschrijven van sleutelfactoren voor deze gedecentraliseerde, gedistribueerde en niet-hiërarchische vormen van samenwerking. Niettemin, de grenzen van dit type productie en het ontstaan van mogelijke organisatorische hybriden moeten nog worden onderzocht.

Deze studie wil een bijdrage leveren aan de literatuur over open source, en peer- en sociale productie door een analyse te maken van de beperkingen en uitvoerbaarheid van deze vormen van productie op het gebied van internetveiligheidsbeleid en de reacties op grote internetveiligheidsincidenten. In zekere zin kijkt deze studie naar twee coole ideële kinderen, peer-productie en internetveiligheid, die elkaar voor het eerst tegenkomen, om te zien hoe ze met elkaar kunnen opschieten.

De eerste, leidende, onderzoeksvraag van deze studie is: Kan het omgaan met bepaalde internetveiligheidsincidenten geclassificeerd worden als peer-productie? Een tweede reeks vragen heeft betrekking op a) de rol en het belang van geheimhouding en de rol en het belang van antagonisten bij reacties op incidenten en b) factoren die ofwel de aanwezigheid of afwezigheid van elementen van peer-productie in tegenreacties ondersteunen.

2. Theoretische grondslagen

Om deze vragen te beantwoorden maakt deze studie gebruik van theorieën, modellen en concepten zoals open source- en peer-productie, het beheer van internet, veiligheidsbeleid, vertrouwen en geheimhouding.

Het idee van peer-productie beschrijft samenwerkingsverbanden voor de ontwikkeling van vrije en open sourcesoftware en Wikipedia. Volgens Benkler wordt deze "nieuwe modaliteit van het organiseren van de productie" vergemakkelijkt door goedkope internet-gebaseerde communicatie, gedeeld eigendom van goedkope ICT-systemen, en daarmee lagere kosten van decentrale samenwerking.

De studie stelt een taxonomie voor om verschillende vormen van sociale productie helder te labelen. *Sociale productie* is de overkoepelende term, die de nieuwe niet-markt-gestuurde en niet-hiërarchische vormen van productie beschrijft. *Peer-productie* vereist dat sociale productie is gebaseerd op een gelijke verdeling van gezag onder degenen die bijdragen leveren. *Commons-based peer-productie* benoemt die deelverzameling van de sociale productie, waarin geen centrale controlerende bestuurseenheid of toe-eigening van geproduceerde goederen bestaat.

De definiërende kenmerken van deze varianten van sociale productie zijn gedistribueerdheid, openheid en socialness. *Gedistribueerdheid* verwijst naar de netwerk topologie van de bijdragende actoren, het ontbreken van centrale of decentrale knooppunten, de prevalentie van peer-bestuur en ad-hoc-hiërarchieën. *Openheid* beschrijft de toegankelijkheid van het productieplatform, haar transparantie, een gedeeld sentiment en de toegankelijkheid en aanpasbaarheid van geproduceerde goederen. *Socialness* verwijst naar het sociale transactie kader, dat wil zeggen, contribuanten nemen niet deel op basis van financiële prikkels of hiërarchische druk, maar op basis van intrinsieke drijfveren.

Door de relaties tussen geheimhouding, veiligheid en sociale productie te ontrafelen, voorziet deze studie in een analytisch model om de aanwezigheid en afwezigheid van elementen van de sociale productie in reacties op veiligheidsincidenten te helpen verklaren. Een zekere mate van geheimhouding gaat samen met peer-productie, maar het verandert de onderliggende economie van het peer-productie model en vermindert de levensvatbaarheid en ideële ambitie van het productiemodel.

3. Onderzoeksopzet

Dit onderzoek maakt gebruik van casestudies van internetveiligheidsincidenten. Dit maakt een uitgebreide beschrijving mogelijk van de actoren die betrokken zijn bij de reacties, de organisatie van hun activiteiten, hun toegang tot relevante informatie, en hun interacties en samenwerking.

Als criteria voor de selectie van de gevallen moeten de incidenten relevant en beperkt in tijd en omvang zijn. Bovendien moeten de gegevens over de incidenten en de respons maatregelen beschikbaar en toegankelijk zijn. Uiteindelijk werden de Estse cyberaanvallen en de Conficker botnet gekozen.

Het identificeren van de toepassing van peer-productie in het reageren op incidenten vereist drie stappen. De eerste taak is het, goederen en diensten, geproduceerd bij reactie activiteiten, te identificeren, wat een beschrijvende verhaal ervan vereist. Ten tweede worden de response activiteiten gecategoriseerd, met behulp van de opgegeven criteria van peer-productie - gedistribueerdheid, openheid, en socialness. De derde opdracht is om te beslissen of de reactie kan worden aangemerkt als peer-productie of niet.

Om te analyseren waarom de elementen van peer-productie wel of niet in de reacties werden toegepast, werd een model van de toepasbaarheid van peer-productie ontwikkeld. Het beschrijft factoren die zijn geïdentificeerd als aannemelijke voorwaarden voor de toepasbaarheid van peer-productie voor het creëren van een bepaald goed.

Belangrijkste bron van empirische gegevens zijn kwalitatieve interviews met experts, aangevuld met bureauresearch en enige directe observatie. Interviews werden getranscribeerd en gecodeerd.

4. Het in gevaar brengen van het internet

De basis voor de verdere analyses werd voorbereid met een historiografische schildering van de twee incidenten.

Gedurende drie weken vanaf 27 april tot 18 mei 2007, werden onderdelen van de Estse internet infrastructuur overweldigd door Distributed Denial of Service (DDoS)-aanvallen, website inbraken, DNS-server aanvallen, massa e-mail en commentaar spam. De aanslagen vormden een keerpunt in de geschiedenis van internetveiligheid vanwege twee aspecten. Ten eerste maakten de aanslagen het voor een breder publiek plausibel, dat cyberaanvallen kunnen worden gebruikt als een instrument in internationale of bilaterale conflicten. Ten tweede vormen de

Estse cyberaanvallen een zeldzaam geval waarin een "aangelegenheid van nationale veiligheid" recht werd gezet door een gemeenschap van technische experts.

In het najaar van 2008 buitte een stukje malware een kritieke kwetsbaarheid uit van het Microsoft Windows-besturingssysteem en installeerde zich snel en geruisloos op miljoenen pc's. De geïnfecteerde computers werden een deel van een enorm botnet, gecreëerd met gebruikmaking van een verbazingwekkende reeks innovatieve aanvalstechnieken. Ondanks zijn ongewone afmetingen, is het botnet alleen gebruikt in kleine gevallen van cybercriminaliteit, waardoor het onderliggende doel tot op de dag van vandaag een mysterie blijft. De zaak Conficker verhoogde de bewustwording van het probleem van botnets onder een breder internationaal publiek. Het liet ook een indrukwekkende wereldwijde samenwerking zien tussen technische experts en werpt licht op de meestal verborgen veiligheidsinstituties van het internet.

5. Het produceren van internetveiligheid

Hoe wordt het internet beveiligd en verdedigd in tijden van crisis, en welke producten en diensten worden geleverd door de reagerende actoren, om de status quo ante uiteindelijk weer te herstellen?

De Estse aanval werd gemitigeerd door de Estse gemeenschap van technische experts, ondersteund door hun internationale collega's. CERT-EE ontwikkelde zich tot het centrale knooppunt voor informatie-uitwisseling en coördineerde een aantal van de defensieve maatregelen. Realisering van veiligheid bestond uit een succesvolle mitigatie van verschillende aanvalstechnieken, in het bijzonder DDoS mitigatie. Situational awareness werd gerealiseerd door het monitoren van technische en sociale netwerken. Samenwerking tussen binnenlandse actoren werd vergemakkelijkt doordat de veiligheid professionals in dit land elkaar kenden. Ad-hoc samenwerking met de wereldwijde veiligheidgemeenschappen moest al improviserend worden opgezet.

In het begin bestond de reactie op de Conficker botnet slechts uit een kleine adhoc groep van veiligheid professionals. De technische aard van de aanvallen, de betrokken internet subsystemen, en de gekozen responsstrategie vereisten een grote respons gemeenschap. De opbouw van de gemeenschap werd mogelijk gemaakt door een half decennium van voorafgaande conferenties en netwerken. De primaire tegenmaatregelen omvatten malware en botnet analyses, defensieve DNS, en het creëren van sinkholes. Reverse engineering van de malware's binaire code was onontbeerlijk om het botnet te begrijpen en de mitigatie te ontwerpen. Door de implementatie van Defensieve DNS, konden bots niet langer contact opnemen met commando servers van het botnet. Sinkholes, databases met informatie over botverkeer, werden gebouwd om kennis te verwerven over de activiteiten en de verspreiding van de bots.

Beide reacties waren afhankelijk van bijdragen van de internetveiligheid gemeenschappen. Hun waarden, normen en interne praktijken beïnvloeden de manier waarop internetveiligheid wordt gerealiseerd na grootschalige incidenten.

6. Sociale dimensies van internetveiligheid

De empirische analyse identificeert de rol van peer-productie in internetveiligheid door uit te puzzelen of de responses op de Estse 2007- en de Conficker incidenten gedistribueerd, open en sociaal waren. In beide gevallen hebben de responses niet voldaan aan alle aspecten van het peer-productie model.

In termen van gedistribueerdheid vermengen beide reacties elementen van decentralisatie en gedistribueerdheid. Zowel CERT EE en de Conficker Cabal hadden een centrale rol in de coördinatie van de reacties. Activiteiten zoals situationele controle, DDoS mitigatie en malware onderzoek werden voornamelijk verdeeld, terwijl Defensive DNS of verkeer filtering werden uitgevoerd op een decentrale manier. Individueel gezag was niet gelijkmatig verdeeld binnen de reactie gemeenschappen. In beide reacties was er grotendeels gelijkwaardigheid tussen de actoren, maar er ontstonden enige interne hiërarchieën. Reactie netwerken konden normen afdwingen onder hun leden.

De reacties in beide gevallen zijn niet open geweest op de manier waarop open source-productie open is. In het beste geval vonden de activiteiten plaats in een omgeving die kan worden omschreven als een gesloten openheid of BarCamps binnen muren. De toegang tot de veiligheidsgemeenschappen is meestal beperkt, vereist doorlichting van- en referenties voor potentiële leden en is gebaseerd op onderling vertrouwen. Echter, de respons netwerken omvatten ook niet-doorgelichte actoren. Binnen de grenzen van de veiligheid gemeenschappen werden vele informatiebronnen gedeeld. Terwijl sommige van de actoren in het verleden hadden geflirt met de ideologie van openheid, is het leidende principe inmiddels verantwoorde openbaarmaking (responsible disclosure) en verantwoordelijkheid ten opzichte van de slachtoffers.

De reacties kwamen het dichtst bij peer-producties in de dimensie van socialness. Drijfveren van contribuanten leken op die, gebruikelijk in open source-projecten, waaronder het doel om idealistische waarden te bevorderen of persoonlijke belangstelling. Het terugdringen van de 'bad guys' en 'spelen met geheimen' kunnen motivaties zijn die uniek zijn in veiligheid gemeenschappen, maar ze passen nog steeds in het open source-motivatie kader. Hetzelfde geldt voor de gedeelde interesse in het verkrijgen van indirecte beheersing. Bovendien hebben zij die een bijdrage leveren aanzienlijke verplichtingen tegenover hun gemeenschappen, en hun loyaliteit aan gemeenschappen kan die aan hun werkgevers overstijgen.

7. Beperkingen van Openheid

Het open source-toegang-voor-allen-beleid en het *swift trust*-model worden vervangen door een systeem van beperkte toegang, het doorlichten van potentiële deelnemers en het delen op een need-to-know basis. Deze uitkomst - geen pure peer-productie, maar significante toepassing van elementen van peer-productie - roept de vraag op waarom bepaalde elementen van peer-productie wel in de reacties kunnen worden gevonden, terwijl anderen niet zijn toegepast.

Dit onderzoek heeft de hindernissen voor openheid geanalyseerd op basis van een model van factoren die een open of gesloten benadering bevorderen. Zij identificeert de antagonist van de gemeenschap, de "bad guys", als de drijvende kracht tot geheimhouding en tot de voorkeur van de internetveiligheid gemeenschappen voor een ommuurde organisatorische aanpak. De aard van de sociale productie in de respons pogingen lijkt op een institutionele aanpak, die probeert om een aantal grote voordelen van het idealtypische peer-productie model in zich op te nemen, terwijl op hetzelfde moment met de behoefte aan geheimhouding rekening wordt gehouden.

De toepassing van groot vertrouwen als eerste vereiste voor het lidmaatschap van de veiligheidsgemeenschappen kan daarbij worden opgevat als een barrière tegen het risico van ontrouw. Het waargenomen hybride karakter van sociale productie vermindert het risico van het binnendringen van kwaadaardige slechteriken, die proberen om verdedigings-inspanningen van de veiligheidsgemeenschappen te niet te doen. Gemeenschap-gebaseerd doorlichten en een zekere mate van toegankelijkheid voor nieuwe deelnemers houden de kosten van geheimhouding relatief laag.

Geheimhouding belemmert één bron van peer-productie effectiviteit: het ongeplande, onbeperkte gebruik van middelen via een groot aantal actoren met verschillende talenten en vaardigheden. Maar de internetveiligheid gemeenschappen kunnen nog steeds profiteren van de relatief goedkope toegankelijkheid voor nieuwe deelnemers om hun voordeel te doen met het verkrijgen van externe informatie.

8. Conclusie

De productie van internetbeveiliging zag er in de onderzochte gevallen anders uit dan gewoonlijk wanneer de openbare veiligheid wordt aangetast. Waar transnationale samenwerking tussen politie, wetshandhaving, of strijdkrachten ontbreekt of ontoereikend is, hebben verspreide bottom-up samenwerking door ad-hoc teams of permanente internetveiligheid gemeenschappen een beslissende rol gespeeld. Anders dan bij open source- en rigide peer-productie, is de toegang tot productie platforms, input bronnen, en intermediaire informatieve producten beperkt, en er is geen cultuur van onvoorwaardelijke openheid en transparantie.

Dit onderzoek heeft een aantal beperkingen. Het geeft slechts een beperkte blik op de relatie tussen peer-productie en internetveiligheid. De observaties laten geen duidelijke conclusies toe over optimale organisatorische ontwerpen van respons inspanningen. Dit geldt nog meer omdat de respons organisaties een element van historische contingentie hadden. Bovendien zijn de werknemer-werkgever relaties niet op intra-organisationele data gebaseerd.

Een aantal hiaten in onderzoek zijn in de loop van deze studie geobserveerd. Internetveiligheid gemeenschappen verdienen nadere analyses vanuit verschillende theoretische invalshoeken en niveaus van analyse - of het nu Ostrom's common-pool gemeenschappen, epistemische gemeenschappen, of Internationale Relaties theorieën zijn. Meer gedetailleerde theorieën over op vertrouwen gebaseerde samenwerking, sociale productie, verspreide, uit netwerken bestaande en egalitaire gemeenschappen zouden waardevol zijn. Tenslotte worden diepere theoretische en studies en studies over het ontwerpen van open beveiligingssystemen aanbevolen.

De studie wordt afgesloten met discussies over de stand van zaken in internetveiligheidsbeleid en ideeën over hoe dat 'open te leggen'. Meer recente trends in internetveiligheidsbeleid hebben de indruk van een relatieve achteruitgang van internetveiligheid gemeenschappen gewekt. Aan de andere kant zou het centraliseren van effecten en hiërarchisering van de internetveiligheid gemeenschap door een reeks van maatregelen kunnen worden vermeden.

Curriculum vitae

Andreas Schmidt has held the position of a doctoral researcher at the Faculty of Technology, Policy and Management at Delft University of Technology. Recent publications encompass papers in established academic peer-review journals such as Telecommunications Policy and International Studies Review, contributions to edited books, as well as numerous presentations at academic conferences. His research has covered the organisation of Internet security production and large-scale incident response, the Estonian cyberattacks, networked governance in international relations, emerging hybrids of networks and hierarchies, and geopolitical aspects of information technology.

Public-oriented outputs include interviews for newspapers as well as appearances in radio and TV talk shows. Schmidt has been a speaker at Internet policy-oriented conferences. Furthermore, he has been an invited panellist, discussant, and participant at events or meetings hosted by the European Commission, the European Parliament, and Bundestag MPs.

In his professional career, he has served as an independent or in IT-focussed business consultancies large enterprise clients since the early noughties. His subject matter expertise comprises project management, process analysis, and the design of operations of high-security environments.

The civil alternative to military conscription allowed Schmidt, born in Bad Arolsen, Germany, on November 7, 1971, to serve as an emergency medical technician. Studies in international relations, philosophy, economics, and computer science in various European cities followed. He holds a *magister artium* degree in political science and history. His 2004 thesis received a *summa cum laude*. Entitled "Hegemony by technology", it explored the ways in which information technology could be instrumentalised for strategic gains in global politics and thereby described the intellectual origins of recently revealed intelligence programs.