RESEARCH-ARTICLE

When User Needs Meet Power: Improving Security Usability by Recognizing Where Business Needs Come First

SIMON EDWARD PARKIN, Delft University of Technology, Delft, Zuid-Holland, Netherlands

Published: 24 August 2025

NSPW '25: New Security Paradigms Workshop
August 24 - 27, 2025
Aerzen, Germany

https://doi.org/10.1145/3774761.3774920
ISBN: 9798400718755

When User Needs Meet Power: Improving Security Usability by Recognizing Where Business Needs Come First

Simon Parkin
Delft University of Technology
Delft, Netherlands
s.e.parkin@tudelft.nl

Abstract

There have been a great number of usability improvements put forward in user security and privacy research. However, it is not guaranteed that beneficial changes proposed in research reach practice. If an improvement is seen not to benefit the service, or to be too difficult or costly to implement, the service owner may ignore it. Equally bad for users is if 'powerful' stakeholders – whoever it is who has the resources and influence to make the usability change in the real world – are selective about which elements of a proposed usability improvement they are willing to implement; this risks diluting the protections that the change would have afforded for users. Here we propose a shorthand 'user second as user-centred' approach to preparing usability improvements to security and privacy technologies and processes. Paradoxically, this perspective promotes usability by prompting a consideration of usability improvements as a value proposition for existing systems, and consideration of how the proposed changes align with stakeholder decision-making criteria. This is as opposed to relying on an assumption of usable security and privacy as being universally beneficial – such an assumption would rely on the powerful stakeholders to appreciate the need for improvement and not dilute it in any way, in the process of transferring it into a real-world service or environment. We show how this approach may be mobilized in an adaptation of the premortem planning technique, and explore a range of case studies where usability needs were variously warped or kept intact, either with the cooperation of powerful stakeholders or without them.

CCS Concepts

• **Security and privacy** → *Usability in security and privacy*; *Economics of security and privacy*.

Keywords

Usable security and privacy, security interventions, security economics, premortems

ACM Reference Format:
Simon Parkin. 2025. When User Needs Meet Power: Improving Security Usability by Recognizing Where Business Needs Come First. In *New Security Paradigms Workshop (NSPW '25), August 24–27, 2025, Aerzen, Germany*. ACM, New York, NY, USA, 13 pages. https://doi.org/10.1145/3774761.3774920

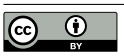

NSPW '25, Aerzen, Germany

ACM ISBN 979-8-4007-1875-5/25/08
https://doi.org/10.1145/3774761.3774920

1 Introduction

The role of the user in security and privacy has received increasing attention in research and practice over time, from academic beginnings as early as the 1970s [63] and with focused attention from the mid-1990s onwards (e.g., [2, 84]). However, the accommodation of human aspects as an intrinsic part of secure systems has not always gone smoothly, with the spirit of usability getting lost in the implementation – this includes password complexity rules which only decades later are accepted as being unfit-for-purpose [8] despite earlier evidence in usability research (e.g., [2]), and workplace behaviours being recognized as important only to be smothered by one-way, one-size-fits-all training [9, 36], to name a few. If an interface or process were to maintain an unusable security or privacy feature after it is evidenced in research, it is assumed to be unpalatable for the maker or manager to leave that feature as it is (such as difficult password policies).

Usability improvements are framed in research as if they are universally-agreed benefits, which any stakeholder should want; any stakeholder who is 'powerful' enough – having enough influence and resource to realize the usability improvement in practice – should invest in bringing each improvement into the real world (be it a large company, or a regulatory entity or government). However, a growing body of research with, for example, security managers in organizations – but also the increasing amount of research around regulatory and compliance needs – has found that a usability improvement may cost a manager or an organization in resources, or impede what they want which does not align with the needs of a usability improvement. For example, targeted security training in a company may be beyond the skill-set, resources, or remit of the security manager [35, 36], or not available from vendors even if companies want it [35].

Up to now, we have been focusing on systems, rather than specific decisions about those systems, and the decision-makers who have power within those systems. This is important given that usability research is increasingly examining real-world interfaces and services, which are managed by someone who makes decisions about those systems. Here we leverage security economics (e.g., [5–7]) and research of sociotechnical systems (e.g., [43, 73]), in combination with lessons from security and privacy usability. We need to understand the decisions, decision preferences, and incentives of the 'powerful' stakeholders that we want to change the minds of.

It could be argued that analysis of decision-making is outside the domain of the usability researcher. However, unless we define the non-divisible form of the improvements we recommend, they risk remaining an indistinct 'nice to have'. This is critical for security and privacy because, presumably, our usability recommendations relate to the management of a particular risk. If there is a narrative

of 'usability helps' rather than 'usability is essential' in security and privacy, the powerful stakeholder is given license to reshape or ignore the intervention. For instance, recommendations to simplify password policies and remove expiration (e.g., [34]) have met with resistance, where in some instances expiration windows are extended but not completely removed.

This becomes a problem because any shortfall in (security and privacy) usability improvements, and in seating them correctly alongside other tasks and activities, leaves residual 'non-usability' that the user still needs to respond to. This happens while the decision-maker has potentially moved on, having regarded the 'user' element as sufficiently addressed *according to their preferences*. The core ask of the paradigm here is, if a researcher expects a specific stakeholder with power to implement the researcher's security/privacy usability improvement, that the researcher should ask themselves, both if and how they really expect this to happen in reality. Two research questions then emerge: RQ1: In what ways can the planning and implementation of security/privacy usability improvements be decoupled from the sole reliance on usability improvements?, and; RQ2: What alternative strategies can be used to plan a strategy for their adoption?

The work described here is around a paradigm that frames security and privacy usability as a value proposition between the user and the maker, with the usability researcher best-placed to do that framing. We provide a summary of the paradigm next. This is followed by an examination of the connections between users and the decision-makers within organizations or services, in Section 2. We then posit one method to mobilize the paradigm, based on the premortem approach (Section 3), followed by case studies in Section 4. We close with Discussion (Section 5) and Conclusions in Section 6.

1.1 Our contribution

The paradigm proposed here is driven by a heuristic of 'user second as user-centred'; assume that there is some criteria that is seen as more important than usability, for the 'powerful' stakeholder who is best-placed to make a usability improvement happen in practice. This could be lack of resources or expertise, or simply that they do not care to implement proposed changes as those changes do not align with their priorities.

This heuristic then shifts the strategy from imploring the stakeholder to put usability improvements first, to positioning a researcher's usability improvements relative to those decision criteria. This decouples the researcher from a dependency on the stakeholder to already 'get it' and be in agreement, and prompts the researcher to explore other ways that the usability improvement can be realized (with – or without – that powerful stakeholder).

This heuristic also rests on a hope that a way can be found to carry a usability improvement into implementation intact and without compromise, rather than it be surrendered to stakeholders and potentially diluted. A stakeholder is unlikely to understand the intentions and 'hard lines not to cross' in how the usability improvement should be implemented, especially as security and privacy usability requires particular skills – ideally the researcher would manage some part of the negotiation. We can perhaps forgive these powerful stakeholders, in cases where the hard requirements of a proposed usability improvement were not clearly defined.

If we anticipate this challenge, it forces us to (i) be steadfast about what is critical in a usability improvement, to protect users from a risk exposed in research, and; (ii) apply the knowledge gleaned from the research, to understand our options for getting the improvement into the world. The latter makes us consider whether we are relying on the powerful stakeholder, or whether there are other ways to reduce the exposure of the user to the identified (usability) risk. This then bears similarities to adoption-centred design [16] and the need for benefits beyond usability – here, that would be risk reduction for the user. There are a range of tenets within this proposal, as below.

- **Usability alone will not activate improvement.** Akin to adoption-centred design (discussed in Section 2), we assume that powerful stakeholders must have their own interests met first, before considering anything else – they may have limited resources, for instance. Up to now, we have assumed that usability is 'equal first' with anything else, as it is implicitly considered to be universally beneficial (and as such, have the capacity in itself to create or release the resources necessary for implementation).
- **Relate usability improvements to the criticality of risks.** There may be some usability improvements which are critical to managing risks for users, beyond 'smoothing the path'. Connecting usability to critical system risks has been considered in previous work (e.g., [9, 34, 57]). Such criticality can be considered alongside the powerful stakeholder's decision criteria, where system risks become a common language.
- **When powerful stakeholders do not act – adversarial usability.** If the powerful stakeholder will not adopt the change we see as necessary to support users, we must decide whether to stop there or drive change another way. There is another side to the perceived universal benefit of usability improvements that we currently hold to – if the powerful stakeholder does not immediately take on the cause, that other stakeholders will implore they do so. Even if this happens, it takes time, where in the meantime the user is left exposed. Borrowing from adversarial interoperability [23], we may look to enact a change even when the powerful stakeholder does not want to cooperate or change their system, or do so in good time – we 'plug in' a more usable alternative (as with E2EE chat apps and password managers, in Section 4).
- **Recognize existing infrastructure.** In considering that security and privacy usability needs may be running second-place to some other concern, this – paradoxically – prompts us to recognize that there is an existing system in place, and existing resourcing decisions and stakeholder incentives. Usability improvements then become a value proposition relative to solutions which are already in place. This extends to resources which are already allocated.
- **Define non-divisible usability improvements, ahead of negotiation.** If the researcher is not able to consider the target ecosystem, they can define what is critical to the usability improvement working, be it the intent or the features.

Negotiation may require softening some expectations to find agreement with others [43]. If the expert researcher is not present to 'escort' the solution, this leaves all the deliberation to the powerful stakeholder who may not appreciate what is needed for the improvement to work. Defining the usability improvement also supports powerful stakeholders to consider softening their own needs, if the value proposition of more usable security provides comparable benefits (e.g., more usable password policies which result in fewer helpdesk calls [57]).

2 Background: Problems and Stakeholders

Here we consider stakeholders who can enact security and privacy usability improvements, how usability relates to business value, and the role of the user relative to these concerns. This informs RQ1.

2.1 Powerful stakeholders in security and privacy

We have referred to 'powerful' stakeholders in the Introduction (Section 1). This refers to stakeholders with influence and resource, who we expect could adopt usability improvements from research and make them happen in a real-world system. Often, we do not consider who this 'powerful stakeholder' would be, again perhaps leaning on an implicit belief that usability improvements are beneficial, and in turn, somewhat inevitable. There may be a more 'powerful' stakeholder than the usability researcher, but this does not mean by default that this stakeholder has the available resources or inclination to enact the usability improvement. If we acknowledge this, it allows us to reframe the dynamic between research and practice, as more akin to 'usability but I need to cover my costs', as a *value proposition*. As an example, self-service helpdesks in organizations rest on an argument of putting the control over workplace credential management into the hands of users, but also allow for organizations to reduce the staffing costs of in-person helpdesks [57].

Research into unintended harms of well-meaning cyber-risk controls [18] refers to a 'service owner', and assumes that this service owner would be interested in reducing unintended harms to their intended users, including harms from burdensome controls. The motivating argument is that the service owner, as the 'powerful' stakeholder, would want to conduct this check; otherwise, users may be irrecoverably impacted, for instance if they are unable to access a managed system to then contact the service owner. However, this checking activity may be both costly and require socio-technical skills which are alien to the decision-maker.

There may be powerful stakeholders who may appropriate usability techniques – such as training and behavioural cues – to impose *more* security upon users, placing security first [35]. These are often security managers such as Chief Information Security Officers (CISOs), whose job is in essence to do exactly that, and put security first. This can be for reasons of regulatory compliance, where secure user behaviour is also increasingly being recognized in regulations. We are then appealing to this powerful stakeholder in the presence of other such priorities.

There are also powerful stakeholders who may impose *less* security and privacy in the presence of other preferences. This could be said to be where, for instance, large social media platforms – and online services in general – are the powerful stakeholder in question. This can result in dilemmas for users such as 'FoMO privacy' (Fear-of-Missing-Out privacy) [79, 80], where users may want to engage with platform content, but begrudgingly provide information that the platform demands in order to enable that engagement, against their own privacy preferences.

2.1.1 *Usability and adversarial interoperability.*

'Big Tech' is one target of Cory Doctorow's arguments, for a need to force change by forcing interoperability of new features into existing, potentially combative digital services and technologies [23]. This is seen as the long route to regaining benefits for users and others in the system, other than the powerful stakeholders who own the technology. 'Indifferent' or 'adversarial' interoperability would be where we do not have the cooperation of the powerful stakeholder. Doctorow also considers a 'cooperative' interoperability where the powerful stakeholder engages with the proposal for change.

In Doctorow's writing, indifferent and adversarial interoperability may be pertinent because the stakeholder does not care for the proposed change, or would actively work to make the proposed change too difficult within the system that they have influence over. If the powerful stakeholder represents a (large) business, they may be fixated on profit-making, in effect prioritizing what has made them powerful and is seen to help maintain that power. This can be at the expense of a more usable experience for users, if the existing experience is regarded by that powerful stakeholder as being 'usable enough' to keep users in the system. This may be when the user is known by the powerful stakeholder to be in a position where they have to engage with the service or product, to get what they want (such as a website offering a distinct product); this facilitates the powerful stakeholder demanding user information and concessions that are unfair or manipulative [1]. This can also be the case for social media platforms [79], where a user is forced to forgo their personal preferences to be able to engage with their contacts (who must all use the same platform in order to talk to each other). This is part of Doctorow's arguments; 'large' – here synonymous with 'powerful' – tech companies can choose to listen to users, advocacy groups, and policymakers at their own discretion, as they cannot be forced to put the user first (despite being user-driven services). This is where regulation may attempt to step in, with policymakers then being another stakeholder to appeal to in such circumstances.

Another perspective which has some parallels to Doctorow's adversarial interoperability has been drawn from contextual interaction theory [12]. Specifically with regards to how policies – or improvements – are implemented, one strategy would be for new actors to 'invite themselves' to negotiations, which in our case may be usability researchers (as experts in an intervention, or as representatives of user needs). Where contextual interaction theory would inform the values and interests of involved stakeholders, these factors may be more directly considered by usability researchers, to assess whether it is feasible to work with or against self-interested stakeholders to improve the situation for users.

2.2 Usability and business value

Powerful stakeholders do not necessarily have limitless resources – they are just the stakeholder we consider to be more likely to have resources than other stakeholders. This could be a large business or association of businesses. We may otherwise regard a stakeholder as powerful, if they have influence over the decision-making of other stakeholders who have resources (such as a regulator). For instance, crafting training to directly meet user needs requires resource to specify, as well as costs in deployment to reach users, and monitoring during maintenance [9]; regulation has driven an adoption of user training in companies [37]. Conversely, usability improvements may be *disruptive* to other concerns that a (powerful) stakeholder has [16], so are not only about benefits.

Put another way, the most powerful stakeholder we can imagine, be it an organization, a regulator, government, or society at large, cannot always conjure up resources and action even if they agree that security and privacy usability would be helpful. For instance, smart device regulations – such as the Cyber Resilience Act (CRA) in the EU – are appearing in the mid-late 2020s to benefit consumers, arguably in response to the Mirai botnet attacks years prior. Governments saw a need to take action, but regulation is not immediate; user-centred features, such as non-trivial device credentials by default, have taken time to regulate, and will take time to mandate and enforce.

The alignment of usability generally with business has been considered in prior work. Prior research in HCI has realized through engagement with stakeholders that "end user benefit alone was not of premiere value to companies" [16], and that there must be benefits beyond usability, even in cases where companies claim that they value the 'user experience'. The work of Chilana et al. [16] also highlights that usability improvements must be "measurable and visible". However, challenges to adoption can be the need to scale to meet end-user needs, but also the very notion of changing existing processes (whether that change is real or expected). This results in a need to consider *adoption-centred design*, reflecting on what the research aims to achieve and inherent design choices.

Similar to the suggestion to consider adoption-centred design, when usability researchers consider business goals, "with the necessary respect of other stakeholders' responsibilities and expertise", this can establish "a common ground for collaboration" between managers and usability experts [38]. This would, however, appear to assume that usable security and privacy complement or add to business needs, and are never disruptive. As an example in security, Florencio & Herley [28] characterized a tension in website security – it was seen as necessary to protect user accounts with passwords, but commercial website owners also saw it as necessary to limit the complexity of their password policies, so as not to be more complicated for users than the policies of competitors' websites.

2.2.1 The alignment of usability with other preferences. Considering how usability improvements may impact the 'managed' nature of a target system, there is a distinction to be made between usability and utility in practice [72]. A proposed change may have *misalignment* with a real-world system, for instance requiring a change in how people already work. A *problematic* usability improvement may indeed realize a benefit to usability, but create drawbacks elsewhere in the system. In this context, here we are considering that usable security and privacy must have utility to the powerful stakeholder, and if it does not, its improvements may not appear as beneficial as they were within the accompanying research. For instance, physical tokens may provide additional security, but research has ultimately concluded that users can be averse to the risk of losing or forgetting the token, and not want to adopt them [3] – online banking and other real-world services have been seen to move toward smartphone apps to address a similar need for convenience, as users are more likely to remember their phone and are already used to looking after it on a daily basis.

Tarkkanen et al. [72] note that new legislation may impose changes to the working of existing processes, where Naqvi et al. [47] note similar issues, that regulatory requirements may be followed steadfastly even when a compliant system is shown to be an unusable system. Earlier work on CISO views of security policies and working practices [57] found that the usability of security technologies can be better considered, if connections to other cost concerns are surfaced (such as reduction of breaches and IT-helpdesk staffing), or if usability relates to regulatory needs, such as user training [37].

There has been prior research that considers user and usability needs against decision-maker preferences, often leveraging security economics principles [9, 34]. Despite the stark framing of user burden in this research, the transfer of usability improvements into practice has often been selective, and decided by the powerful stakeholder, who decides how the proposal for usability fits into the picture after their own priorities are preserved. Research in the area of security and privacy usability has, in essence, aimed to frame usability in a similar way to core business concerns such as finance and risk, to motivate powerful stakeholders to prioritize usability. Regardless of how successful this may have been, it has not ensured that usability techniques are translated intact, to benefit the user first.

2.3 User involvement and power structures

Usability improvements would ideally realize a benefit to concerns other than usability, as alluded to by Tarkannen et al. [72] in their consideration of the utility of usability improvements. Within security and privacy, it would seem reasonable to argue that a usability improvement should realize a reduction of critical risks, or reduction in workarounds or other behaviours which undo risk-management measures (such as when users write down passwords in a way that makes them easy for others to see). Where Nielsen's usability criteria differentiate between minor and major usability problems [49] (where the latter may induce errors), we may need to consistently articulate criticality of the risks addressed by usability proposals, in a similar fashion. The urgency of a usability improvement contributes to managing risks for users, as a common language – or exchangeable information within a *trading zone* between stakeholders [70] – for many stakeholders in the conversation.

Where users are impacted by a lack of usability, it may be that it is only a relative few among many users who are severely impacted [55], even though many may be silently struggling with unusable security and privacy controls as *hidden costs* [57]. If end-users of

public or commercial services, such as apps, experience these burdens, many may complain at one focal place, namely through app marketplace reviews [55]. Across organizations, even if users in different companies are using – and struggling with – the same enterprise security solutions, there is no means currently for *collective bargaining* [61] for usability improvements beyond the boundaries of an individual organization (e.g., employees arguing for being paid by their company, for the time needed for their work computers to start up every day [51]). This diminishes the power of workers to collectively argue about burdens in the population-wide manner framed by Herley [34]. This informs whether a researcher examining the usability of a particular security and privacy control would take on the stakeholder of one organization, or a wider community of organizations all using the same control (foremost by pitching their improvements to regulators or governments).

2.3.1 'Power to' and 'power over' stakeholders. In representing burdened, often non-expert users, researchers in security and privacy may be assuming a 'power to' dynamic with the powerful stakeholders responsible for systems they research [65], i.e., *a power to* request a usability improvement. If usability improvements were a universally-agreed benefit, we might expect that the powerful stakeholder must eventually implement the improvement, conversely because they should have no argument for not implementing it (if usability is universally agreed to be good).

Cybersecurity and privacy research does not outwardly assume 'power over' stakeholders [65] such as government, regulators, service owners, or security managers, but that it appeals to concerns that they share. On the face of it, usability improvements should not be optional or up for debate for the user as they add more protection, and yet when it comes to implementation, they somehow are. Business needs override the improvements, lack of business resources means no action is taken, and if security teams are themselves under-resourced then risks are dumped with users. Usability proposals would benefit from relating to those business needs and limitations, otherwise it inadvertently assumes 'power over' those concerns.

In framing usability as universally beneficial, and especially if the associated risk reduction is not communicated, usable security and privacy research can appear aspirational. This is highlighted by Turnhout et al. as an issue with much co-production research between multiple stakeholders [73] (for instance in environmental sustainability, in their case). There can be unequal power relations, with outcomes governed by elite actors (which here we may regard as companies, manufacturers, governments, regulators, and the like). This puts pressure on researchers to be clear about the minimal working version of their usability improvement, so that it does not become diluted. Arguably, much behavioural research has been co-opted as a way to bombard users with generic security training (as a form of user support), manipulative interfaces (rather than persuasive design [31]), and predatory privacy features (where users must give up data as a means to move beyond the uncertainty and complexity of reaching their goal [1]).

If hard lines are not drawn, parties may need to soften their aims during negotiation in order to reach consensus [43]. We might then ask if usability improvements can define more explicit hard lines. For instance, where an early study in user security explored allowing users ten attempts to get a password correct rather than three [13], we might accept that ten attempts increases the risk of compromise by a determined attacker, but draw a line, say, at five attempts to still realize a user benefit. In an environment under regulatory oversight, it also becomes beneficial to the powerful stakeholder if the requirements of a usability improvement align with the regulations or laws which, for instance, dictate that stakeholder's permission to access a market. However, user needs are somewhat under-specified in areas such as training & awareness, for example (that training must be provided, but less is said about what to do if it does not work for a user [36, 37]). Public or open services may need to provide universal access, as in the European Union [52], and account for accessibility needs [59]. Any regulatory developments would ideally shift usability costs to the 'least cost avoider', as the party who would incur the lowest cost to avoid the identified harm [45].

3 From paradigm to pathway

In this section we address RQ2, exploring alternative approaches. We propose a means to embody the principles of the 'user second as user-centred' paradigm for usability assessment of security and privacy technologies. This approach is based on Klein's *premortem* method [40, 42]. When a project plan is being assessed, a premortem exercise assumes that the project has already failed, and asks contributors to "generate plausible reasons for the project's failure", as a way to identify risks to success early on. The approach is informed by *prospective hindsight*, as imagining that an event has already occurred. This is argued to prevent contributors from fixating on plans or goals which may need to be replaced.

Within security and privacy usability, it can be argued that much research relies on powerful stakeholders seeing the value in a proposed improvement, as the one and only 'plan'. Premortems have been proposed elsewhere to, for instance, anticipate cybersecurity incident response planning [25]. We envisage a premortem serving as an additional step after a usability assessment, as detailed in the steps described in Section 3.1 and shown in Figure 1). This shows where the plan toward a 'user-centred' goal could be reached in a different way, and the energy behind the research utilized effectively.

3.1 A candidate approach

The following steps describe one candidate approach for anticipating how to align a security or privacy usability improvement with the powerful stakeholder(s) who are best-placed to support it. Other possible approaches are discussed in Section 3.3.

Step 1. Usability assessment. Conduct a usability assessment and identify changes to interfaces or processes, which will benefit the user. This may identify interface or process problems, unmet or unrecognized user needs, etc.

The usability assessment is assumed to be carried out by an independent usability researcher, such as a researcher at an academic institution separate from the service or environment being studied. However, this arrangement is increasingly meeting with challenges, given that usable security and privacy research is more often examining digital products and services used in the real world,

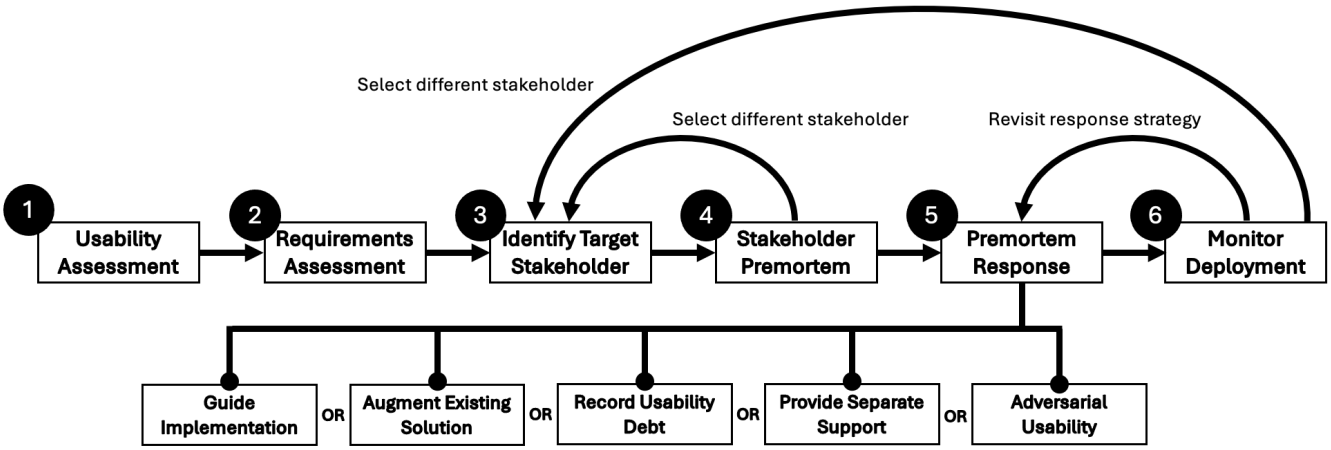

Figure 1: The steps in the method for positioning usability improvements into managed systems.

i.e., closed, commercial products (such as device operating systems and updates [78], smart home products [75], mobile apps [27], etc.). Academic researchers may have independence, but this may conversely be a barrier to access and influence. Another issue with academic researchers being outside of the product or service they are assessing is that their efforts may be less easily communicated to the powerful stakeholder (as has been seen between HCI research and policy [69]), unless some common remit and terminology is established [21, 70].

Alternatively, the usability assessment may be conducted by someone internal to an organization or who has been hired. For instance, a designer who has to shape client requests into an interface design may have some limited capacity to raise usability issues. There are examples of this, where designers refuse a request to integrate a 'dark pattern' technique into a website that has been commissioned [83], in essence positing a less-unusable alternative that is not so fixed to benefiting the powerful stakeholder.

A usability assessment may also be driven in a way that user needs are gathered by asking technical experts what they think users need [37, 48, 75]. We would only regard this as a usability assessment if the expert has directly engaged with representative users, and is able to comment on user needs objectively without concern for upsetting a superior or client (i.e., without already deciding that user needs are going to come second to business needs, in effect before even seeing what the user needs are). However, user-facing advice and policy has been seen to rely on the input of technical experts rather than users themselves [48].

Step 2. Establish requirements of the usability improvement. Ideally, articulate the (cyber-)risk that would affect the user if the change is not enacted (the residual risk), and its criticality. We may not be able to say what the impact would be like on a target system, but we should aim to describe aspects of the deployment which resonate with stakeholders. Also, articulate the 'hard lines' of the usability proposal, as anything beyond that may need to be negotiated in order to reach consensus. This all relies on the researcher, in their capacity in 'escorting' the proposal to powerful stakeholders. An example may be where user success and failure rates are presented for different combinations and complexities of authentication to a service, alongside user throughput, etc.

Step 3. Identify the target stakeholder. Identify the stakeholder who needs to enact the change. Look ahead to consider what is in it for them. Consider what it is about them that means that they must make the change, and are able to achieve it. This may be, for instance, because the stakeholder is optimally placed, regarded as naturally sympathetic to the user need, or because regulatory expectations can be interpreted as requiring them to adopt the usability improvement (as with much research with, e.g., the security of medical records). It could also be that the stakeholder is the owner of the service, as is generally assumed when CISOs in organizations are looked to for improvements to employee security.

Step 4. Stakeholder decision-making premortem. Consider the needs and pressures of the target stakeholder. For instance, many studies have focused on the pressures on CISOs in organizations [35, 62]. It may be somewhat hopeful – even naive – to identify a stakeholder in advance, but a premortem approach allows for considering failure to reach that stakeholder, and alternative plans, as in the next step.

As a second part of this step, assume that the chosen stakeholder does not choose to implement the usability improvement. Consider reasons why they would not want to do it. If a reason is that they will not listen to a researcher, but may listen to another (powerful) stakeholder, reconsider whether that is the target stakeholder, and return to Step 3 above. For instance, businesses may listen to their sector-specific regulators, as with how security training is increasingly present within organizations [36, 37]. This can potentially be done as a group exercise with stakeholders of varying 'power', especially if the capability to enact the improvement is spread across multiple parties – there may then be a shared incentive structure. For example, the security of home smart devices may rely on manufacturers, but also regulators, retailers, and Internet

Service Providers (ISPs) [29, 75].

Step 5. Respond to the premortem forecast. It would be necessary to choose a strategy for how to respond to the needs of stakeholders, as identified in the previous step – options include:

- **Guide non-divisible implementation.** if the stakeholder is cooperative, work with them to implement the change, escorting the change in a manner that watches for any compromise of what is needed relative to Step 2 (as in the final step in the process, Step 6).
- **Incorporate the change into an existing process of change.** This may be more realistic if the usability improvement is clearly articulated as being part of a regulatory change, for instance. Consider also that a system may be replaced as a part of infrastructure renewal, allowing for a new and more usable replacement to be put in place – there may be an opportune moment to implement the change, by augmenting an existing solution. Some negotiation of the usability improvement, as more 'realistic' expectations, may happen here.
- **A record of usability debt.** This can be used in conversations with the stakeholder, or by that stakeholder in future tendering and procurement processes, as a means to gather resources later which they do not have now (e.g., for procuring employee-facing security solutions [14, 36]). Minimal requirements can then be used as ultimatums to suppliers, especially if the stakeholder knows that compromising those minimal requirements may only perpetuate historical usability problems.
- **Provide support to users if stakeholder is not able to.** A stakeholder may lack resources or the will to enact the change, or believe they have already done enough for their service/system users. This could already be considered 'adversarial' if it works 'outside' of the functionality and influence of the service. For example, governments may provide (independent) advice for how to stay safe on social media platforms; if the stakeholder is sympathetic but lacks resources, they may choose to signpost the support resource.
- **Adversarial usability, if stakeholder is not responsive.** If the stakeholder is uncooperative for whatever reason (be they 'indifferent' or 'adversarial'), it may be necessary to pursue a process of making the usability improvement happen some other way, without their involvement (and potentially with their opposition). This is also where the goal of the usability improvement should be meaningful even if detached from the target stakeholder (and potentially the infrastructure they have power over). One example of where detachment seemed to be considered necessary was the 'adversarial' movement to provide more secure, encrypted chat apps (which popularized the likes of Signal) in the wake of the Snowden revelations.
  This part of the premortem exercise decouples success as being 'whatever the stakeholder can implement' to preserving the user need and exploring other ways to achieve it. This may require returning to Step 3, to identify a new stakeholder to consider engaging with. This is to accept that the user is not the stakeholder's first priority, prompting a revisit of the strategy to meet user needs.
  To consider this as a different perspective on existing approaches, an example would be additional controls on social media platforms – prior research has provided mock-up interfaces to mimic popular platforms for, e.g., visual encryption features [26]. This primarily serves ecological validity, but could equally be re-framed as a proposition to the mimicked platform, of 'work with us to adopt this, or we will find another way to implement it for users'.

Step 6. Monitor deployment. Observe the deployed change or other action, to capture if it becomes diluted (and why), or fails to get off the ground. This aligns closely with the causal analysis of why behavioural interventions can fail in practice [54]. The premortem approach lends itself to the shift advocated by Osman et al., from a question of 'what went wrong?' to 'what could go wrong, and how could it go wrong?'. Similar to the stakeholder premortem, the outcomes of the monitoring activity may require engagement with a different stakeholder (Step 3), or revisiting the premortem response (Step 5). In the context of a premortem, advance consideration of what to observe for success can inform what measurements need to be in place, and what a failure condition would look like. For instance, despite the proliferation of encrypted chat apps, many years into their availability there are users who do not trust them as being secure [20].

3.2 Timelines & timing

A usability improvement may have more success if it coincides with a renewal of the infrastructure that the powerful stakeholder cares about (as seen in the 'Augment Existing Solution' path, in Figure 1). For instance, if an organization's legacy accounts-management infrastructure is being replaced, the stakeholder may be looking for ways to update or future-proof the new systems for a long time – this would be an opportune time to revisit the usability of services such as authentication mechanisms, and update the user-facing solutions. It could be argued that user-facing policies in an organization are more likely to be changed when IT systems need replacing, or when there is a new CISO wishing to put their stamp on the security apparatus, rather than any regular, ongoing process of adjusting policies to better fit users. Put bluntly, usability improvements are just as likely (if not more) to arise as part of a planned IT renewal, rather than because of usability problems reported by the users.

Considering time in another way, the time required to move a usability improvement into practice must also be considered, and would ideally match in some way to the timelines of the powerful stakeholder. Certainly, a real-world organization is unlikely to stop everything to let researchers try out a new usability solution; just as with test-beds for vulnerability patching in organizations, it may be possible to trial an improvement with a subset of users (thereby also evidencing its benefits). If a usability improvement were to take a long time, it might not match with the timelines of the powerful stakeholder; this harks back to Anderson's characterization of OS provision as 'ship it on Tuesday and get it right by version 3' [5], and how commercial interests may require getting the product out

to gain market share, satisfying the needs of any other stakeholder (including the user) once that has been achieved. Usability in regulations and development would then be an ex-ante approach, and threat modeling when a product is nearing or at deployment an ex-post approach.

3.3 Other approaches

Other approaches are possible, aside from the premortem approach described here. Usability may be an afterthought in product development processes [15]; Shostack has explored a consideration of user perspectives in threat modeling for digital products, as a means to account for the power dynamics that favour the view of (more technical) developers [66]. This was inspired by the work of Slupska & Tanczer [68] to provide a threat modeling approach to confront developers with the misuse potential of their products (as extreme consequences of a lack of usability, as threats to personal safety). Other work by Sim et al. [67] provides a method which prompts security and privacy designers to recognize discernible groups of users who may use their product.

There then may be ways to consider the usability needs of users within existing design and development processes. The approach described in Section 3.1 would be one which 'plugs into' these processes. We could also consider a Backcasting approach [82], starting with a desired state of both users and businesses getting what they want or would need, and working back from there – this would similarly require an understanding of what all parties need, but would raise questions as to where and how negotiation would happen.

4 Case Studies

Here we present a series of case studies, examining prior usability developments through the lens of the 'user second as user-centred' paradigm. For each of these case studies, we list the kinds of usability improvements that researchers have been finding, and how far they go in shaping the change, relative to the steps of the process listed in Section 3. We accept that these case studies are retrospective – we know what happened and what the outcome was. The aim here is not to convey some predictive capabilities of the paradigm, but instead to recast a range of prior efforts as stepping away from a reliance on the powerful stakeholder seeing the value of the usability improvement, all while users are exposed to an identified risk that is exacerbated by a lack of usability.

4.1 Dark patterns

Some online services and websites have – knowingly or otherwise – used interface design to go beyond persuasive design and into manipulative design, creating *dark (or deceptive) patterns* [31]. Examples include creating urgency with sales countdown timers that run on a loop, or baseless/hard-coded warnings of low item stock. Regulations are gradually catching up with these dark patterns, such as GDPR in Europe, and the California Consumer Privacy Act (CCPA). However, regulations take time, and cases brought by authorities can only target a few of the offenders at a time to levy fines against.

Usability improvements here currently focus on the implementation of 'fairer' patterns [58], which could be argued are presently defined as values or principles [19], or as a call for information-gathering activities that bring the stakeholder and user together, toward balancing their preferences (at least toward a 'neutral' interface that does not manipulate the user). However, businesses may prioritize website designs which benefit customer throughput and retention, and the collection of customer data, even when they are challenged on pursuing a dark pattern [83].

Researchers have in the meantime created 'tip line' websites (or a 'hall of shame'), for example the Dark Patterns Tip Line[1]. In terms of researchers escorting the improvement, the name-and-shame websites categorize websites according to specific kinds of dark pattern (as may be identified in research). If a powerful stakeholder – in this case, the business that runs an offending website – does not want to engage with the research, the tipline has named the website anyway, and indicates where they are misleading users. This approach arguably chips away at the problem in much the same way as regulations, but is more agile. This then is a kind of adversarial usability, with the goal of influencing manipulative online services to change their practices. Referring to Step 5 as in Figure 1, this also serves as an open *record of usability debt.*

4.2 Small-to-Medium Business (SME) security

There is much government policy and academic research that would call on small-to-medium businesses to invest more in cybersecurity. Although this problem is not 'usability' in the strict sense of interface or process design, it speaks to the user difficulty in obtaining more usable solutions in the first place [56]. SMEs may agree that security is important, but often cannot make resources suddenly appear for the purposes of improving their security – many have limited resources, diverse IT, and may not have a dedicated IT team, let alone a dedicated person responsible for IT-security [22].

The usability call here is to adjust expectations for SMEs, or to provide assistance to make cybersecurity improvements more affordable or accessible, relative to the resources and non-uniformity of smaller business IT [56]. At present, policymakers and the security community seemingly rely on an organization wanting to avoid the future regret of a breach and having not invested in cybersecurity in the present. This in itself does not make IT infrastructure appear. It may, however, not be realistic to expect security to be provided completely for free either.

Governments may want small businesses to be more secure. The 'adversarial' arrangement is then between SMEs and, in a way, policy advisors and the security community (the kind of elite actors mentioned in Section 2.3, on the payroll to provide 'expert' advice). SMEs are the stakeholder who should strive for more security, but concedes that they are not able to achieve it, or questions why it should be them taking action when the voices calling for more security are other actors [53]. One approach which is not necessarily adversarial to SMEs is then to provide support from another stakeholder. In this case, an option for support is the centralization of security guidance and incident recovery to cyber-insurance providers [46]. Arguably, this still requires SMEs to pay for security, but this has to be offset against a model which would have SMEs each invest in creating this infrastructure for themselves.

[1] https://darkpatternstipline.org/

This is an example where the simplest solution would have been to keep responsibilizing SMEs to invest more in security [53], but although they are not dismissive, the solution did not fit. In Figure 1, this would be a response to indifference by powerful stakeholders in Step 5.

4.3 Password managers

The emergence of the password manager market may or may not have been knowingly 'adversarial'. However, it addressed a problem inherent in password systems which otherwise would have been addressed, foremost, by urging users to develop better memory for remembering numerous complex passwords; this was arguably an impossible goal. To reduce the burden on users would have also required service owners to coordinate, and to collectively appreciate the user cost that they were each adding to, and then act collectively to reduce it. It was 'cheap' for service owners to have password protection on their websites, because they each did not bear a cost when a user creates a new credential and adds it to their existing set (only considering the cost if it impeded their competitiveness [28]). Facebook Connect and Google log-in have appeared over the years, and the FIDO Alliance is in a way attempting to address the problem by creating a unifying credential for multiple services that a user may be using. However, initiatives such as FIDO/FIDO2, and Passkeys still face adoption challenges [44], so the usability is some time away – we have not replaced passwords yet.

Sidestepping these problems, password managers and browser password stores emerged as a means to reduce the reliance on users to change their password habits. Password manager applications emerged and took hold because we never made passwords really work well out in the real world in the first place, all despite a long history of trying – and doing so by appealing to the user to do more. It was only later that the original advice was accepted as being unfit for the expanded use of passwords [8], and in the meantime password managers addressed the problem. This is another case of action that improves usability in the face of indifference from existing stakeholders in the ecosystem, though in this instance by providing a collective response to inaction by many, more powerful, stakeholders.

4.4 Tailored training

There are many studies which advocate targeted security training for employees (e.g., [11, 60]). There are other works which recommend to revisit the design of important security measures, to reduce the effort required for employees to use them (e.g., [9, 36]). However, when these recommendations have been communicated out to powerful stakeholders such as CISOs and training vendors [37], what we see is a market of general training which aims to address all security needs in a general way [36, 37] – it is beneficial to awareness managers to deploy a lot of training, for themselves to look busy and warrant their role [36]; CISOs who manage awareness as a minor role are not experts in human behaviour [57]. User needs are being 'met' by deploying interventions which are assumed by their volume to meet all possible needs, without the costly upfront identification of specific needs, or crafting of targeted training. The latter may be an unrealistic expectation in practice, as it would require all organizations to craft the most usable training for their workforce(s) individually; perhaps what has saved organizations up to now is the generic nature of IT in most office-based workplaces.

One response appears to be to recast the role of training managers as 'listening' to employees' security needs [36], as an 'inverse usability' which fixes problems if and when they arise. This, however, does not fix technologies which were not usable in the first place, even if it does identify where employees are having problems; this approach assumes that technology works and addresses problems for anyone who struggles to use it. This achieves the same end of arriving at usable technologies, as advocated in human-centred research [9]. However, it carries a user burden specifically for those (supposedly few) users who may have particular difficulty using solutions provided to them by the organization. Researchers may then ask whether training approaches, as they currently stand, have made sufficient progress or need more attention. Relating to the method illustrated in Figure 1, here, the awareness practitioners who practice a listening approach, are not so much acting against usability, but finding an alternative way to achieve the same shared end of achieving secure working for the majority of users. However, currently this is facilitated by a lack of definition in standards, for the role or expectations of an awareness manager – awareness managers are filling the vacuum with engagement activities that connect to employees and their needs [36].

4.5 Security culture

Parallel to training, there is also 'security culture' as a growing trend, where one version of this co-opts research directions which recognize human aspects of security in organizations. There is a tilt to convince employees to increasingly value security, but often without messaging that directly links security to the core organization culture. This means a risk of having two 'cultures' being cultivated in the one organization. The concept of a distinct 'security' culture is arguably on a path where it is needing to somehow, in the end, supersede the core organizational culture. This only intensifies the tension between working securely and getting work done, if those two 'cultures' are not aligned. This also perpetuates the position of the user as 'small' and having to accept the burdens placed on them. This serves as an example of where a targeted improvement (introducing a 'security culture' to improve the enactment of security behaviours [60]) has to recognize that there is already a 'culture' in almost any organization, as 'the way things are done around here' [11]. It may, however, not be realistic to expect security managers to directly connect with organization culture within their remit; there may though be scope to, for example, connect security to the core values of the business, to reduce confusion among employees as to how security tasks connect to their work [39].

Governments have provided adjacent training and advice, and there are free training resources of varying quality by a range of providers. There is not a lack of solutions, but the researchers' role in guiding what training to use, how to articulate user needs, how to target training to those needs, and do so within the resources and capabilities of an organization, would benefit from researchers continuing to 'escort' the solution into practice. This could, for instance, involve relating behaviour change efforts to core business values, rather than trying to impose security as a core value, thereby inching closer to supporting the context of use. This would be an

approach of working in the midst of a community of adversarial stakeholders, where security managers are increasingly required to provide awareness and training content to employees, yet lack the resources and skills to engage with employees [37]. These concerns relate to the intention of Klein's more recent, extended premortem method, the 'double-barreled premortem' [41]. This prompts managers to consider the risks of inaction, and of not taking risks to try to improve the existing situation. Here, the current situation would be to leave users to navigate competing 'cultures'.

4.6 Tech-abuse support

Spyware can be difficult to spot on a smartphone. Even being able to assess whether there is spyware on a phone can require a high level of technical skills. There is then a need for tech-abuse support to honour a user's *belief* that a personal smartphone has been compromised, and at the same time provide practical support. Frontline tech-abuse support services lack sufficient technical skills to keep up with evolving threats, and often need to acquire those skills from outside [71].

A team of researchers created a spyware clinic [33], for smartphone users to bring their device and have it checked for spyware. This can fill a gap where the 'powerful' stakeholder is not capable of running the service themselves, but also informs future procurement decisions, being run as a kind of 'prototype'. This would demonstrate the costs and outcomes, to inform questions about resourcing. It does, however, become a different question to determine how frontline support charities could keep such capabilities running over time, and whether a centralized service would be possible while also providing access for users. Havron et al. noted that there were a few clinic cases where it was not possible even for their experts to confirm a lack of spyware [33] – it would remain key to ensure that users do not need to have advanced technical skills in order to be heard. Efforts such as the spyware clinic inform Requirements Assessment (Step 2, as in Section 3.1 and Figure 1.

5 Discussion

Here we discuss wider considerations around the paradigm (Section 1.1) and how it might be implemented (Section 3).

5.1 An extended usability assessment for security and privacy

The paradigm in this paper aims to decouple usability improvements from a reliance on the 'powerful' stakeholder in the target ecosystem, to implement some version of the intended change. We can explore what the improvement looks like when it remains intact, and what else we can do if the powerful stakeholder takes no action. For usability improvements in general, we might argue that implementing any aspect of the improvement *is* a success, as it improves usability. However, in usable security and privacy, we must draw hard lines as to what should be implemented to remove a security or privacy risk for the user(s) – doing it halfway will not achieve the intended result. There is a risk that some elements of an intended usability approach may be taken and others left out, or usability elements used instead to sustain the power of the stakeholder, e.g., providing privacy control choices, but only ones which leave the service's access to profitable user data intact.

Prior work combining human factors with security economics [34] has argued about the size of the user burden relative to what it would achieve for security. Other work leveraging security economics has argued about reducing the burden of individual security tasks to use less of a user's 'budget' for security [9]. Here, we consider how stakeholder preferences may distort a proposal for a usability improvement, such that any reduction in user burden is considered only second to what is within the means and interests of more 'powerful' stakeholders than the researcher. These stakeholders arguably have less appreciation than the researcher, for why the improvement had to have the form it did when proposed, to realize the intended improvement for users. Referring back to those earlier works, password policies still often include forced expiry [34], and training became the way to support users rather than (more expensive) re-design of technologies to be more usable [9] – stakeholder interests led the change.

Prompting researchers to frame their usability improvements as a value proposition prepares us to negotiate for those improvements, seeing that allowing them to be reshaped by other stakeholders can lose the intended improvement or create harms of its own [18]. For usable security and privacy improvements, we aim to address risks critical to the user – this would shift the framing of usability improvements from a 'nice to have' to addressing severe risks for users. Usability premortems could help researchers to 'pick which fight' to take on, or more so, which route will keep the improvement intact, subject to negotiation of the value proposition.

5.2 Recasting the image of the user

It may be argued that the 'expert' user knows full well how to use what is already available (noted as early as Whitten & Tygar's work on encrypted communications [81], and loaded visual metaphors as cues). Making an interface or process more usable may require additional investment upfront to understand user needs [9, 64], or would expose that user-facing processes are much less efficient once user effort is recognized as not having a zero-cost [34]. The easiest direction of travel for powerful stakeholders has so far been to see users being told to become security experts; much research has followed this direction.

Similar to comments elsewhere [24], researchers should consider ways to elevate the role of the user in negotiations about security. For instance in security training, it is unlikely that 'the user' is referring to high-ranking executives with the power to push back on onerous demands for user effort – when research of this nature refers to 'the user', it is a rank-and-file employee. The expert guiding the intention of the research, the security community as 'powerful' stakeholder, has decided the place of the user. Yet, framing usability issues as collective, society-wide issues depicts bigger, community-level and societal costs, as with Herley's calculations of collective user effort [34].

Within this consideration of a greater collective of users, it of course is possible that the users themselves carry their own usability needs forward into engagements with more powerful stakeholders. An underlying assumption here is that one (not-as-powerful) user acting alone would face great difficulties, in both being heard and negotiating a usable solution. Yet, ideally, evidence of security and privacy usability challenges should speak for itself, regardless of

who wields it. Usability researchers can translate findings to different usability challenges; a user can convey their own experience in their specific context of use, and articulate their own goals within a process of negotiation, as the exchangeable artifact [70].

It could be argued that researchers carrying the usability improvement forward on behalf of users only serve to perpetuate the image of users being incapable of secure behaviour [24], in order to fit with the pre-existing views of the 'powerful' elite actors who may be the ones resourced enough to make any change happen (especially as some may be paid to be present in the conversation as their professional role) [73]. We may look to re-cast the user's goal as having the same validity as 'decision-maker' studies in organizations or risk management have in studies of 'managers', for instance, which focus on preserving and serving their decisions, rather than (broadly speaking) changing their behaviours to accept burden. There have been prior cases, for instance, of user communities navigating takeover of smart device product lines after the manufacturer stopped support or went out of business [32, 50]. We may ask ourselves how much we know about what users are wanting to do – and do securely – compared to what we know about the needs of powerful stakeholders. This may also help to elevate the needs of users who are failed by user-facing interventions, where Osman et al. [54] note for behavioural interventions generally, that "a recurring theme across many studies and types of failures is that subgroups matter".

Alongside representing society-wide usability issues, as in the work of Herley [34], we can then also go in the other direction, and define hard lines for more specific, smaller user groups. Existing research has already considered discernible groups of users across a range of intersections, be it interpersonal risks [71], at-risk and marginalized groups [77], and accessibility needs [59, 76]. Research with at-risk users already advocates to engage with knowledgeable stakeholders early [10], but we may consider this in the context of adversarial interoperability [23], and where user needs are so critical that something must be done with or without the powerful stakeholder (so that related usability research does not also contribute to keeping the user 'small' by reinforcing, e.g., user responsibilization [17]). In their proposal for human-centred threat modeling, Usman & Zappala [74] propose to identify risks for specific user groups, but also to identify barriers to adoption of protective measures. Where their mechanism of 'reflection' encourages researchers to go beyond their own assessment and to also capture the perceptions of the user, this could – for instance – be extended to bring human-centred risk management into negotiations with powerful stakeholders (the same stakeholders who may be able to alleviate barriers to protective measures). Ultimately, the user is not being asked directly to define what is *reasonable for them*, or being supported to do so. In most cases, the market decides it around them – they are provided for, but not catered to.

5.3 Future work

Future work can further explore how to define and delineate the rigid and negotiable elements of a usability improvement. To further inform our understanding of who we are negotiating with on the path to adoption, we can also engage archetypal 'powerful' stakeholders in qualitative research (for instance, semi-structured interviews), such as security managers in organizations, and policymakers. The aim here would be to further understand how they distinguish one user from another in terms of security and privacy needs, and how they determine whether security provisions are a good fit for those users (e.g., where policymakers rely on expert views [48]).

Work around digital security and privacy often borrows from the world of safety. Discussions in safety are moving beyond a focus on 'human error'. Security and privacy research may similarly move away from responsibilizing the end-user [17], toward ensuring that the environment around them is secure. Future work would then also explore where lessons can be transferred from the safety realm to cybersecurity, especially in terms of ensuring that users are not left to carry residual (usability-related) risks.

Relating to the timelines for interventions (Section 3.2), one throttle on the pace of change would be that an intervention in one environment is not necessarily transferable to another. By more closely relating an intervention to users' context-of-use, it becomes more specific to that context. Connecting this to a need to negotiate, further research will explore a range of heuristics which can serve as a 'toolkit' for stakeholder engagement. For example, we may act to ensure that expected user costs for avoiding a cyber-risk are lower than the costs of having been otherwise subjected to that risk [34]; we may also manage the infrastructure around users so that they are not going beyond their capacity [9].

5.4 Limitations

It may be argued that the extended role we are speculating about here for usability researchers is not appropriate for their remit. However, the researcher can articulate what risks the usability improvement is protecting against, and crucially, the 'minimal' – or non-divisible – specification of the usability improvement, for it to survive exposure to corresponding business decisions and alleviate user risks. This would ideally help to limit the negative externalities perpetuated for users [34] even when some diluted form of the proposal is implemented.

We assume that the usability improvements needing attention are those which would be easy for a powerful stakeholder to reshape or dismiss, as a moral hazard [30] that creates risks for others and no consequence for them. Some usability improvements in security and privacy may genuinely be 'light' in some way; a problem arises when light and severe risks of not improving usability look the same from outside the research. This anticipates researchers lobbying for the importance of their usability improvements; one ready way to do that is if it relates to a regulatory need. Alas, another is if the risk is existential or relates to a real threat of physical or psychological harm [4], as in tech-abuse research [33, 71].

Here we have implicitly been considering the process as being potentially 'asynchronous' – the researcher would prepare for negotiation without necessarily being in direct consultation (yet) with the stakeholder who may ultimately help to realize a proposed usability improvement. This has a limitation that the 'business case' may not be a fit in practice to the concerns of the target stakeholder, if developed without them, but the hope is that it would provide an artifact for shared discussion.

6 Conclusion

Researchers in usable security and privacy would ideally consider how their proposals fit into the real world, to address user needs. We must not assume that 'powerful' stakeholders are willing or able to implement any usability improvement that we propose in research. We have to consider plans which do not rely on usability being the top priority for those stakeholders.

The aim of the work in this paper is to encourage security and privacy researchers to think about how to align a proposed usability improvement with existing priorities in the real world. If we do not expect usability improvements to be implemented, we can also plan for how to fill gaps in user support, if the technology will not be made easier to use. In answering our first research question (RQ1), we might otherwise consult another – more willing and able – stakeholder, to make sure usability improvements are implemented. We might otherwise go against the interests of powerful stakeholders who do not want to make the usability improvement a reality, treating those stakeholders as adversarial to user needs. In answering our second research question (RQ2), the premortem method described here aims to embody this change of mindset for the usability researcher. The approach highlights where the usability expert is well-placed, to *negotiate* for the needs of the user. This process is inspired by a range of approaches, such as adoption-centred design and decision-making in security economics.

Acknowledgments

Thanks to Karen Renaud (NSPW pre-shepherd) for guidance before the workshop, attendees of NSPW 2025 for their comments on the work, and Alexandra Dirksen (NSPW post-shepherd) for input on preparing the final version of the paper. Thanks also to Jenny Lieu and Wijnand Veeneman for many previous conversations which informed parts of this work, and to Yi Ting Chua for providing comments on an earlier version of this paper.

References

[1] Alessandro Acquisti, Curtis Taylor, and Liad Wagman. 2016. The economics of privacy. *Journal of economic Literature* 54, 2 (2016), 442–492.
[2] Anne Adams and Martina Angela Sasse. 1999. Users are not the enemy. *Commun. ACM* 42, 12 (1999), 40–46.
[3] Seb Aebischer, Claudio Dettoni, Graeme Jenkinson, Katarzyna Kinga Krol, David Llewellyn-Jones, Toshiyuki Masui, and Frank Stajano. 2017. Pico in the wild: Replacing passwords, one site at a time. (2017).
[4] Ioannis Agrafiotis, Jason RC Nurse, Michael Goldsmith, Sadie Creese, and David Upton. 2018. A taxonomy of cyber-harms: Defining the impacts of cyber-attacks and understanding how they propagate. *Journal of Cybersecurity* 4, 1 (2018), tyy006.
[5] Ross Anderson. 2001. Why information security is hard-an economic perspective. In *Seventeenth annual computer security applications conference*. IEEE, 358–365.
[6] Ross Anderson and Tyler Moore. 2006. The economics of information security. *science* 314, 5799 (2006), 610–613.
[7] Ross Anderson and Tyler Moore. 2007. Information security economics–and beyond. In *Annual international cryptology conference*. Springer, 68–91.
[8] BBC News. 2017. Password guru regrets past advice. https://www.bbc.com/news/technology-40875534.
[9] Adam Beautement, M Angela Sasse, and Mike Wonham. 2008. The compliance budget: managing security behaviour in organisations. In *Proceedings of the 2008 new security paradigms workshop*. 47–58.
[10] Rosanna Bellini, Emily Tseng, Noel Warford, Alaa Daffalla, Tara Matthews, Sunny Consolvo, Jill Palzkill Woelfer, Patrick Gage Kelley, Michelle L Mazurek, Dana Cuomo, et al. 2024. SoK: Safer digital-safety research involving at-risk users. In *2024 IEEE Symposium on Security and Privacy (SP)*. IEEE, 635–654.
[11] John M Blythe, Lynne Coventry, and Linda Little. 2015. Unpacking security policy compliance: The motivators and barriers of employees' security behaviors. In *Eleventh symposium on usable privacy and security (SOUPS 2015)*. 103–122.
[12] Hans Bressers and Cheryl de Boer. 2013. Contextual interaction theory for assessing water governance, policy and knowledge transfer. In *Water governance, policy and knowledge transfer*. Routledge, 36–54.
[13] Sacha Brostoff and M Angela Sasse. 2003. "Ten strikes and you're out": Increasing the number of login attempts can improve password usability. (2003).
[14] Lina Brunken, Annalina Buckmann, Jonas Hielscher, and M Angela Sasse. 2023. "To Do This Properly, You Need More Resources": The Hidden Costs of Introducing Simulated Phishing Campaigns. In *32nd USENIX Security Symposium (USENIX Security 23)*. 4105–4122.
[15] Deanna D Caputo, Shari Lawrence Pfleeger, M Angela Sasse, Paul Ammann, Jeff Offutt, and Lin Deng. 2016. Barriers to usable security? Three organizational case studies. *IEEE Security & Privacy* 14, 5 (2016), 22–32.
[16] Parmit K Chilana, Amy J Ko, and Jacob Wobbrock. 2015. From user-centered to adoption-centered design: a case study of an HCI research innovation becoming a product. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*. 1749–1758.
[17] Partha Das Chowdhury, Karen Renaud, and Ingrid Ott. 2025. Would 'Secure' Users Lead to Secure Commons? Surprisingly Not!: A framework to evaluate effective power and collective outcomes in cybersecurity.. In *New Security Paradigms Workshop 2025 (NSPW 2025)*. Association for Computing Machinery (ACM).
[18] Yi Ting Chua, Simon Parkin, Matthew Edwards, Daniela Oliveira, Stefan Schiffner, Gareth Tyson, and Alice Hutchings. 2019. Identifying unintended harms of cybersecurity countermeasures. In *2019 APWG Symposium on Electronic Crime Research (eCrime)*. IEEE, 1–15.
[19] Tim de Jonge, Hanna Schraffenberger, Jorrit Geels, Jaap-Henk Hoepman, Marie-Sophie Simon, and Frederik Zuiderveen Borgesius. 2025. If Deceptive Patterns are the problem, are Fair Patterns the solution?. In *Proceedings of the 2025 ACM Conference on Fairness, Accountability, and Transparency*. 3131–3137.
[20] Sergej Dechand, Alena Naiakshina, Anastasia Danilova, and Matthew Smith. 2019. In encryption we don't trust: The effect of end-to-end encryption to the masses on user perception. In *2019 IEEE European Symposium on Security and Privacy (EuroS&P)*. IEEE, 401–415.
[21] Albesë Demjaha, David Pym, and Tristan Caulfield. 2021. Found in translation: co-design for security modelling. In *International Workshop on Socio-Technical Aspects in Security*. Springer, 108–128.
[22] Vassilis Dimopoulos, Steven Furnell, Murray E Jennex, and Ioannis Kritharas. 2004. Approaches to IT Security in Small and Medium Enterprises.. In *AISM*. Citeseer, 73–82.
[23] Cory Doctorow. 2019. Interoperability: Fix the Internet, Not the Tech Companies. https://www.eff.org/deeplinks/2019/07/interoperability-fix-internet-not-tech-companies
[24] Andrew C Dwyer, Clare Stevens, Lilly Pijnenburg Muller, Myriam Dunn Cavelty, Lizzie Coles-Kemp, and Pip Thornton. 2022. What can a critical cybersecurity do? *International Political Sociology* 16, 3 (2022), olac013.
[25] Josiah Dykstra, Jamie Met, Nicole Backert, Rebecca Mattie, and Douglas Hough. 2022. Action bias and the two most dangerous words in cybersecurity incident response: An argument for more measured incident response. *IEEE Security & Privacy* 20, 3 (2022), 102–106.
[26] Sascha Fahl, Marian Harbach, Thomas Muders, Matthew Smith, and Uwe Sander. 2012. Helping Johnny 2.0 to encrypt his Facebook conversations. In *Proceedings of the eighth symposium on usable privacy and security*. 1–17.
[27] Adrienne Porter Felt, Elizabeth Ha, Serge Egelman, Ariel Haney, Erika Chin, and David Wagner. 2012. Android permissions: User attention, comprehension, and behavior. In *Proceedings of the eighth symposium on usable privacy and security*. 1–14.
[28] Dinei Florêncio and Cormac Herley. 2010. Where do security policies come from?. In *Proceedings of the Sixth Symposium on Usable Privacy and Security*. 1–14.
[29] Vaibhav Garg. 2021. Covenants without the sword: Market incentives for cybersecurity investment. In *TPRC49: The 49th Research Conference on Communication, Information and Internet Policy*.
[30] Lawrence A Gordon, Martin P Loeb, and Tashfeen Sohail. 2003. A framework for using insurance for cyber-risk management. *Commun. ACM* 46, 3 (2003), 81–85.
[31] Colin M Gray, Yubo Kou, Bryan Battles, Joseph Hoggatt, and Austin L Toombs. 2018. The dark (patterns) side of UX design. In *Proceedings of the 2018 CHI conference on human factors in computing systems*. 1–14.
[32] Gareth Halfacree. 2023. Insteon's Smart Home Cloud Closes Down, But the Open Source Home Assistant Leaps to the Rescue. https://www.hackster.io/news/insteon-s-smart-home-cloud-closes-down-but-the-open-source-home-assistant-leaps-to-the-rescue-e4dbd26c78ce.
[33] Sam Havron, Diana Freed, Rahul Chatterjee, Damon McCoy, Nicola Dell, and Thomas Ristenpart. 2019. Clinical computer security for victims of intimate partner violence. In *28th USENIX Security Symposium (USENIX Security 19)*. 105–122.
[34] Cormac Herley. 2009. So long, and no thanks for the externalities: the rational rejection of security advice by users. In *Proceedings of the 2009 workshop on New security paradigms workshop*. 133–144.
[35] Jonas Hielscher, Uta Menges, Simon Parkin, Annette Kluge, and M Angela Sasse. 2023. "Employees Who Don't Accept the Time Security Takes Are Not Aware

Enough": The CISO View of Human-Centred Security. In *32nd USENIX Security Symposium (USENIX Security 23)*. 2311–2328.
[36] Jonas Hielscher and Simon Parkin. 2024. " What Keeps People Secure is That They Met The Security Team": Deconstructing Drivers And Goals of Organizational Security Awareness. In *33rd USENIX Security Symposium (USENIX Security 24)*. 3295–3312.
[37] Jonas Hielscher, Markus Schöps, Jens Opdenbusch, Felix Reichmann, Marco Gutfleisch, Karola Marky, and Simon Parkin. 2024. Selling Satisfaction: A Qualitative Analysis of Cybersecurity Awareness Vendors' Promises. In *Proceedings of the 2024 on ACM SIGSAC Conference on Computer and Communications Security*. 2666–2680.
[38] Kasper Hornbæk and Erik Frøkjær. 2008. Making use of business goals in usability evaluation: an experiment with novice evaluators. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. 903–912.
[39] Iacovos Kirlappos, Adam Beautement, and M Angela Sasse. 2013. "Comply or Die" Is Dead: Long live security-aware principal agents. In *Financial Cryptography and Data Security (FC 2013)*. Springer, 70–82.
[40] Gary Klein. 2007. Performing a project premortem. *Harvard business review* 85, 9 (2007), 18–19.
[41] Gary Klein. 2025. Double-Barreled Pre-Mortems: Balancing downside risks with upside risks using the popular pre-mortem method. https://www.psychologytoday.com/us/blog/seeing-what-others-dont/202507/double-barreled-pre-mortems.
[42] Gary A Klein. 2011. *Streetlights and shadows: Searching for the keys to adaptive decision making*. Mit Press.
[43] Johannes Franciscus Maria Koppenjan and Erik-Hans Klijn. 2004. *Managing uncertainties in networks: a network approach to problem solving and decision making*. Psychology Press.
[44] Leona Lassak, Elleen Pan, Blase Ur, and Maximilian Golla. 2024. Why Aren't We Using Passkeys? Obstacles Companies Face Deploying FIDO2 Passwordless Authentication. In *33rd USENIX Security Symposium (USENIX Security 24)*. 7231–7248.
[45] Tyler Moore and Ross Anderson. 2011. Economics and internet security: A survey of recent analytical, empirical, and behavioral research. (2011).
[46] Gareth Mott, Sarah Turner, Jason RC Nurse, Jamie MacColl, James Sullivan, Anna Cartwright, and Edward Cartwright. 2023. Between a rock and a hard (ening) place: Cyber insurance in the ransomware era. *Computers & Security* 128 (2023), 103162.
[47] Bilal Naqvi and Kari Smolander. 2024. Practitioners' Perspectives on and Prospects for Usable Security. *Computer* 57, 10 (2024), 66–74.
[48] Lorenzo Neil, Harshini Sri Ramulu, Yasemin Acar, and Bradley Reaves. 2023. Who comes up with this stuff? interviewing authors to understand how they produce security advice. In *Nineteenth Symposium on Usable Privacy and Security (SOUPS 2023)*. 283–299.
[49] Jakob Nielsen. 1992. Finding usability problems through heuristic evaluation. In *Proceedings of the SIGCHI conference on Human factors in computing systems*. 373–380.
[50] NL Times. 2023. Alternative app that can unlock VanMoof bikes more popular amid bankruptcy fears. https://nltimes.nl/2023/07/15/alternative-app-can-unlock-vanmoof-bikes-popular-amid-bankruptcy-fears.
[51] NL Times / ANP. 2024. Supreme Court rules that obliged early arrival for work counts as paid hours. https://nltimes.nl/2024/09/13/supreme-court-rules-obliged-early-arrival-work-counts-paid-hours
[52] Official Journal of the European Union. 2023. European Declaration on Digital Rights and Principles for the Digital Decade. https://digital-strategy.ec.europa.eu/en/policies/digital-principles.
[53] Emma Osborn. 2015. Business versus technology: Sources of the perceived lack of cyber security in SMEs. (2015).
[54] Magda Osman, Scott McLachlan, Norman Fenton, Martin Neil, Ragnar Löfstedt, and Björn Meder. 2020. Learning from behavioural changes that fail. *Trends in Cognitive Sciences* 24, 12 (2020), 969–980.
[55] Simon Parkin and Yi Ting Chua. 2022. A cyber-risk framework for coordination of the prevention and preservation of behaviours. *Journal of Computer Security* 30, 3 (2022), 327–356.
[56] Simon Parkin, Andrew Fielder, and Alex Ashby. 2016. Pragmatic security: modelling it security management responsibilities for SME archetypes. In *Proceedings of the 8th ACM CCS International Workshop on Managing Insider Security Threats*. 69–80.
[57] Simon Parkin, Aad Van Moorsel, Philip Inglesant, and M Angela Sasse. 2010. A stealth approach to usable security: helping IT security managers to identify workable security solutions. In *Proceedings of the 2010 New Security Paradigms Workshop*. 33–50.
[58] Marie Potel-Saville and Mathilde Da Rocha. 2023. From dark patterns to fair patterns? Usable taxonomy to contribute solving the issue with countermeasures. In *Annual Privacy Forum*. Springer, 145–165.
[59] Karen Renaud. 2021. Accessible cyber security: the next frontier?. In *Proceedings of the 7th International Conference on Information Systems Security and Privacy (ICISSP 2021)*. 9–18.
[60] Karen Renaud and Wendy Goucher. 2014. The curious incidence of security breaches by knowledgeable employees and the pivotal role a of security culture. In *Human Aspects of Information Security, Privacy, and Trust (HAS 2014)*. Springer, 361–372.
[61] Jake Rosenfeld. 2021. You're paid what you're worth: And other myths of the modern economy. In *You're Paid What You're Worth*. Harvard University Press.
[62] Kimberly Ruth, Veronica A Rivera, Gautam Akiwate, Aurore Fass, Patrick Gage Kelley, Kurt Thomas, and Zakir Durumeric. 2025. "Perfect is the Enemy of Good": The CISO's Role in Enterprise Security as a Business Enabler. (2025).
[63] Jerome H Saltzer and Michael D Schroeder. 1975. The protection of information in computer systems. *Proc. IEEE* 63, 9 (1975), 1278–1308.
[64] M Angela Sasse and Ivan Flechais. 2005. Usable security: Why do we need it? How do we get it? O'Reilly.
[65] Hanna Schneider, Malin Eiband, Daniel Ullrich, and Andreas Butz. 2018. Empowerment in HCI-A survey and framework. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. 1–14.
[66] Adam Shostack. 2025. Who Are "We"? Power Centers in Threat Modeling. Rossfest Symposium in memory of Ross Anderson.
[67] Mattea Sim, Kurt Hugenberg, Tadayoshi Kohno, and Franziska Roesner. 2023. A scalable inclusive security intervention to center marginalized & vulnerable populations in security & privacy design. In *Proceedings of the 2023 New Security Paradigms Workshop*. 102–115.
[68] Julia Slupska and Leonie Maria Tanczer. 2021. Threat modeling intimate partner violence: Tech abuse as a cybersecurity challenge in the internet of things. In *The Emerald international handbook of technology-facilitated violence and abuse*. Emerald Publishing Limited, 663–688.
[69] Anne Spaa, Abigail Durrant, Chris Elsden, and John Vines. 2019. Understanding the Boundaries between Policymaking and HCI. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. 1–15.
[70] Jonathan M Spring, Tyler Moore, and David Pym. 2017. Practicing a science of security: A philosophy of science perspective. In *Proceedings of the 2017 New Security Paradigms Workshop*. 1–18.
[71] Leonie Maria Tanczer, Isabel López-Neira, and Simon Parkin. 2021. 'I feel like we're really behind the game': perspectives of the United Kingdom's intimate partner violence support sector on the rise of technology-facilitated abuse. *Journal of gender-based violence* 5, 3 (2021), 431–450.
[72] Kimmo Tarkkanen, Ville Harkke, and Pekka Reijonen. 2015. Are we testing utility? Analysis of usability problem types. In *Design, User Experience, and Usability: Design Discourse: 4th International Conference, DUXU 2015, Held as Part of HCI International 2015, Los Angeles, CA, USA, August 2–7, 2015, Proceedings, Part I*. Springer, 269–280.
[73] Esther Turnhout, Tamara Metze, Carina Wyborn, Nicole Klenk, and Elena Louder. 2020. The politics of co-production: participation, power, and transformation. *Current opinion in environmental sustainability* 42 (2020), 15–21.
[74] Warda Usman and Daniel Zappala. 2024. SoK: A Framework and Guide for Human-Centered Threat Modeling in Security and Privacy Research. In *2025 IEEE Symposium on Security and Privacy (SP)*. IEEE Computer Society, 33–33.
[75] Veerle Van Harten, Carlos Hernandez Ganan, Michel Van Eeten, and Simon Parkin. 2025. "All Sorts of Other Reasons to Do It": Explaining the Persistence of Sub-optimal IoT Security Advice. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*. 1–19.
[76] Yang Wang. 2017. The third wave? Inclusive privacy and security. In *Proceedings of the 2017 new security paradigms workshop*. 122–130.
[77] Noel Warford, Tara Matthews, Kaitlyn Yang, Omer Akgul, Sunny Consolvo, Patrick Gage Kelley, Nathan Malkin, Michelle L Mazurek, Manya Sleeper, and Kurt Thomas. 2022. Sok: A framework for unifying at-risk user research. In *2022 IEEE Symposium on Security and Privacy (SP)*. IEEE, 2344–2360.
[78] Rick Wash, Emilee Rader, Kami Vaniea, and Michelle Rizor. 2014. Out of the loop: How automated software updates cause unintended security consequences. In *10th Symposium On Usable Privacy and Security (SOUPS 2014)*. 89–104.
[79] Fiona Westin and Sonia Chiasson. 2019. Opt out of privacy or" go home" understanding reluctant privacy behaviours through the FoMO-centric design paradigm. In *Proceedings of the New Security Paradigms Workshop*. 57–67.
[80] Fiona Westin and Sonia Chiasson. 2021. "It's so difficult to sever that connection": The role of FoMO in users' reluctant privacy behaviours. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. 1–15.
[81] Alma Whitten and J Doug Tygar. 1999. Why Johnny Can't Encrypt: A Usability Evaluation of PGP 5.0.. In *USENIX security symposium*, Vol. 348. 169–184.
[82] Wikipedia. 2025. Backcasting. https://en.wikipedia.org/wiki/Backcasting.
[83] Leah Zhang-Kennedy, Maxwell Keleher, and Michaela Valiquette. 2024. Navigating the Gray: Design Practitioners' Perceptions Toward the Implementation of Privacy Dark Patterns. *Proceedings of the ACM on human-computer interaction* 8, CSCW1 (2024), 1–26.
[84] Mary Ellen Zurko and Richard T Simon. 1996. User-centered security. In *Proceedings of the 1996 workshop on New security paradigms*. 27–33.