Bias and noise in security risk assessments, an empirical study on the information position and confidence of security professionals

de Wit, Johan; Pieters, Wolter; van Gelder, Pieter

Important note
To cite this publication, please use the final published version (if applicable).
Please check the document version above.

ORIGINAL ARTICLE

Bias and noise in security risk assessments, an empirical study on the information position and confidence of security professionals

Johan de Wit[1] · **Wolter Pieters**[2] · **Pieter van Gelder**[1]

Abstract

Professionals working in both the physical and cybersecurity domain need to assess and evaluate security risks. As information on risks in general and security risks in particular is often imperfect and intractable, these professionals are facing a challenge in judging both likelihood and consequences, but how much do their existing psychological biases play a role in these judgments? In this paper, we present new empirical evidence on the perception of the information position and confidence levels of security professionals, the influence of detailed information and the conjunction fallacy, and the level of noise in security assessments. This paper adds to the literature by examining, for the first time, risk assessments by professionals in realistic, real life, security cases. The results show clear indications for overconfidence, comparative ignorance, influence of the conjunction fallacy, and influence of individual experience on security decision making in the professional security domain. The observed phenomena might have far reaching effects on security risk management in organizations and society.

Keywords Security assessment · Confidence · Information position · Conjunction fallacy · Decision biases · Noise

JEL Classification D810

✉ Johan de Wit
johan.de.wit@siemens.com

1 Delft Technical University, Faculty of Technology, Policy and Management, Building 31, Jaffalaan 5, 2628 BX Delft, The Netherlands

2 Behavioural Science Institute, Radboud University, Thomas van Aquinostraat 4, 6525 GD Nijmegen, The Netherlands

Introduction

The security risk field is dealing with malicious, and therefore, man-made, risks. These risks vary from physical security risks like intrusions, theft, holdups all the way to cyber security risks like hacking attempts, ransomware attacks, and IP theft. Nowadays these two domains converge as physical and cyber attacks and threats collide into hybrid threats. To manage these risks, both governments and organizations have introduced security management processes and security staff to assess, evaluate, and manage security risks (ANSI/ASIS 2012; ASIS_International 2015). Security staff, further referred to as security professionals, are educated and trained to perform these tasks. They need to decide, on a daily basis, which risks to take into account, decide how to evaluate them and which security controls to implement.

These decisions are not easy though. In the case of future events originating from complex interactions between multiple independent human agents, occurrence frequency or probability data are often lacking. The assessment of the uncertainty of security risks, therefore, is often based on expert judgment rather than based on evidence or objective data (Möller 2012; Talbot and Jakeman 2011).

As part of their role security professionals are expected to address this uncertainty and form a predictive judgment. Their judgment is often the primary input for risk decisions and allocation of resources (Alruwaii and Brooks 2008). At the same time, human decision making has proven to be not only based on reasoning but is prone to mental short cuts or heuristics, and biases which are defined as systematic deviations from reasoning (Gigerenzer and Selten 2002; Kahneman 2012; Simon 1982; Slovic 2000; Tversky and Kahneman 1975). As the security of society and organizations is thus heavily depending on the individual, subjective judgment of security professionals, understanding their decisions based on their assessments, is paramount to understand security risk management.

In this paper, we present the results of a study in which we ask security professionals to indicate their information position (the level of availability of precise information and/or evidence) when assessing security risks, and to estimate the likelihood of realistic security events for which we vary the descriptions to explore the influence of more or less information. These experiments are based on the conjunction fallacy, predicting that likelihood estimates increase when case descriptions have more specific information, whereas they should actually decrease. Beside the corresponding bias in security risk judgments, the predictive judgments of the individual security professionals might show noise, i.e., a between-subject variance in likelihood estimates within a single condition, where one would hope that different experts give similar judgments instead.

This empirical study will answer the following questions:

- Do security professionals usually have exact information on security risks,
- Are they usually confident about their predictive judgments,
- Would more information grow their confidence,
- Is their judgment of likelihood depending on more or less information,
- Do security likelihood judgments vary under influence of the conjunction fallacy.

The influence of individual expertise of these questions is analyzed. As the future cannot be certain by nature, professionals might be expected to 'know that they cannot possibly know' (known unknowns). Based on this the confidence of security professionals in their predictive judgments can be expected to be limited.

In the next section the theory on security risks, predictive judgments, expert judgment, bias, and noise are briefly discussed. In the section research method the experiments and survey setup are detailed followed by a section in which the results are analyzed. The paper ends with a discussion section and conclusions.

Theory and background

Security teams are tasked to manage security risks to keep them at an acceptable level. The individuals responsible for managing and accessing security risks, in this study referred to as security professionals, often, if not always, apply a risk management process of some sort to structure their assessment.

Risks are defined as the effect of uncertainty on objectives (ISO 2018). In this definition an effect is understood as a deviation, positive or negative, from the expected, often referred to as consequences. The uncertainty of risks is usually referred in terms of their likelihood. "Uncertainty is the state, even partial, of deficiency of information related to, understanding or knowledge of, an event, its consequence, or likelihood" (ISO 2019, p. 6). The understanding and judgment of a risk are, thus, related to the availability of information about it. As Hansson states: 'Knowledge about a risk is knowledge about the unknown' (Hansson 2012, p. 34).

Various risk management processes consist of subsequent process stages: establishing the context, risk identification, risk analysis, risk evaluation and risk treatment (ANSI/ASIS/RIMS 2015; Information_Security_Forum 2018; ISO 2018; ISO/IEC 2011). They also stipulate feed-back loops to establish an on-going, recurring process. As explained in the introduction, this is inherently a decision-making activity, involving decisions on how to evaluate and treat the risks.

Entering the domain of decision making opens up centuries of research, debate, and established theories and practices. Individual decision making is studied ever since the ancient Greek philosophers. As Aristotle stated: the origin of action is choice, and that of choice is desire and reasoning … good action and its opposite cannot exist without a combination of intellect and character' (Allingham 2002). During the last half century renown scholars have unraveled human decision making and especially the cognitive processes guiding them (Baron 2004; Carbone et al. 2017; Slovic 2010).

So far, however, little scientific studies are conducted exploring individual decision making by security professionals in their daily praxis of security risk decision making. These professionals play a decisive or advisory role in security risk treatment; hence, they are determining or at least influencing the security in organizations and society. Understanding their individual preferences and priorities, and the role of information and uncertainty, is of vital importance to understand their security risk judgment.

Judgment of the uncertainty component of risks is related to the deficiency, or in other words availability, of information of an event. Intractable uncertainty is the result of a lack of information than cannot possibly be known (Kahneman et al. 2021). Even with unlimited resources and/or time this information cannot possibly be learned. On the other hand there is imperfect information, information that could be known but is not. Risk decision makers can decide to retrieve more information and enhance their imperfect information position. Often decision makers should or could know that the information they need to decide on is imperfect or even intractable. Many decision makers, however, seem to ignore their lack of information. This attitude is referred to as objective ignorance (Kahneman et al. 2021). The obvious fact that the future is hard or even impossible to predict is often ignored by decision makers (Jain et al. 2013). This attitude of ignorance allows decision makers to have confidence in their decision making, and they mistake their confidence for predictive validity (Kahneman et al. 2021).

In the security domain, both intractability and imperfect information contribute to a lack of risk information and a situation of ambiguity, a situation in which likelihoods either do not exist or are not known (Carbone et al. 2017). It is, therefore, often supplemented or even replaced by subjective expert judgment (Möller 2012).

Expert judgment is considered a degree of belief, based on tacit knowledge and expertise (Cooke 1991). Subjective interpretation, further referred to as judgment, forms the primary input for security risk assessments and risk management processes. Individual judgment is based on the available information, tacit knowledge and 'hard-to-measure' expertise. As this judgment is meant to assess risks, which are possible future events, it is referred to as predictive judgment. The outcome of some of these predictive judgments might become clear in the (near) future and in this cases these judgments can be verified. Examples of these are weather forecasts or predictions on elections. If the predictive judgments involve probabilistic predictions they are often, if not always, non-verifiable (Kahneman et al. 2021). If for example the predictive probabilistic judgment of a risk materializing is 15%, whether or not this particular risk materializes does not allow to verify the judgment. The probability judgment of 15% means this risk materializes 15% of the times in similar circumstances. This prediction of 15% will be valid whether or not this risk materializes. Only after a substantial amount of time and 'similar circumstances,' it might become clear if 15% of the time in similar circumstances is a valid predictive judgment. Due to characteristics of security risks and their large variety of modus operandi, the similarity of circumstances is questionable and thus predictions for security risks can be regarded as non-verifiable by nature.

The huge body of knowledge on judgment and decision making under risk has identified numerous flaws in individual assessments and judgment. Beside biases, which are defined as systematic deviation, human judgment is susceptible to noise (see Fig. 1). Previous work by the authors concluded that security professionals are vulnerable to decision biases to the same extent as lay people (de Wit et al. 2021). Noise, or precision, is the unwanted variability in professional individual judgments. When confronted with the exact same context and information individuals, even trained professionals, can reach different conclusions, often even very different based on personal characteristics (Andersson et al. 2020). Noise or system noise can

Fig. 1 Target shooting as metaphor explaining bias (accuracy) and noise (precision), reprinted with permission from "Noise: How to Overcome the High, Hidden Cost of Inconsistent Decision Making" by Daniel Kahneman, Andrew M. Rosenfield, Linnea Gandhi, Tom Blaser. Harvard Business Review, October 2016. Copyright 2026 by Harvard Business Publishing; all rights reserved

be differentiated in between subjects noise: level noise, and within subjects noise: pattern noise and occasional noise (Kahneman et al. 2021). Level noise is a categorial, systematic, difference between individuals. Based on personal beliefs, convictions or opinion the judgment of one individual can systematically differentiate from the judgment of another individual (Andersson et al. 2020). A security professional can for example be more risk averse in general than another and based on that reach other judgments. Pattern noise is an individual, case by case, variation of an individual. Some specific aspects of security risks can evoke a stronger response by a security professional for example because of previous experiences (Dumm et al. 2020). So the judgment of an individual professional on average might show high risk tolerance except for, for example, holdups where this individual can be very risk averse due to a personal experience. Finally there is substantial evidence that noise is influenced by the occasion. The time of day, the weather, mood etc. influences judgment of individuals.

The influence of the phenomena bias and noise on human judgment has led many scholars to question the viability of such uncertainty assessments. Still in many domains, like security, there are no alternatives or objective procedures available (Hansson 2012; Möller 2012; Tversky and Kahneman 2004). Therefore, predictive, intuitive judgments of uncertainty play an essential role in these decisions (Charness et al. 2020; Kuhn and Sniezek 1996; Tversky and Koehler 1994).

In this study for the first time, to the best of our knowledge, security risk assessments by security professionals are analyzed to explore the influence of information on bias and noise. The respondents in this study are confronted with case descriptions of realistic security risk assessments and are asked to assess the level of likelihood of each case. By randomly varying the presented information between groups of subjects variations of the likelihood assessments can be observed. These variations might be caused by both biases (accuracy) and noise (precision). Comparing the average group assessments shows possible biases (between group comparison) while the within group analysis shows possible noise.

A convenience sample of practitioners form both the security and cybersecurity domain are confronted with realistic security cases with a varying level of information to explore the influence of more or less detailed information on individual

likelihood assessments. These experiments relate to the renowned conjunction fallacy. This fallacy identifies a phenomenon that shows that more detailed information of a situation leads humans to perceive an event as more likely. Logic reasoning, however, would lead to the exact opposite conclusion. Various other scholars have identified very consistent behavior influenced by the conjunction fallacy (Bonini et al. 2004; Fantino et al. 1997; Fiedler 1988; Gigerenzer 1991; Hertwig and Gigerenzer 1999; Ludwin-Peery et al. 2020; Stolarz-Fantino et al. 2003; Tentori et al. 2004; Tentori and Crupi 2012; Tversky and Kahneman 1983).

Many of these studies, however, are based on hypothetical situations in laboratory settings which do not seem to explain real-life behavior (Charness et al. 2020). These studies often involve lay people as respondents who might not be representative for real-life decision makers as risk taking is domain specific (Charness et al. 2020). Our study, on the other hand, investigates judgments of security practitioners on realistic, real-life, cases. The experiments in this study compare between subjects judgments based on different sets of information. The conjunction fallacy is very suitable to explore the systematic deviation caused by more or less detailed information.

In this study several phenomena regarding information, judgment and confidence are explored in the professional security domain. First professionals working in the security domain are questioned about their information position when assessing likelihood and consequences of security risks in real life. As risks are uncertain by nature and especially on risks in the security domain information is often limited or lacking, it is expected that security professionals will acknowledge this. Second: based on this expected meager information position it is hypothesized that security professionals might show modest confidence in their assessments. Third: more experience, training, and education, thus, building individual expertise, on the other hand, is expected to raise and individuals confidence level. Fourth: the possible differences in individual likelihood assessments (noise) are inquired. It is hypothesized that professionals with comparable expertise will reach comparable likelihood assessments in identical case studies. Finally it is expected that varying detailed security case information, by applying the conjunction fallacy, will influence likelihood assessments of security professionals.

Research method

For this study an online survey is set up with Qualtrics survey software. We will investigate both the physical security as well as the cyber security domains. However, related, the physical and cybersecurity domain differ in risk and threat context. The surveys for the two domains are kept identical except for the case descriptions of the two cases as will be detailed below.

The survey starts with questions on the information position of the security professionals in real life on both likelihood and consequence, the two main components of a risk assessment. They are asked how often they:

- Know the likelihood exactly,

- Do not know the likelihood exactly but have quantified information,
- Do not know the likelihood exactly but can estimate the likelihood,
- Do not know the likelihood exactly and cannot estimate it.

The respondents can answer these questions using a five point Likert scale: always, most of the time, about half of the time, sometimes, never. These four questions are repeated for the consequences. The results of these questions indicate the real-life information position of the security professionals in this study and might confirm the position of many scholars that in (security) risk assessments often accurate information is lacking.

These questions are followed by questions about the confidence the respondents feel about their assessments for both likelihood and consequence. A third question asks if the respondents would feel more confident if they would have more information about security risks. The respondents can answer these questions using a similar five point Likert scale: always, most of the time, about half of the time, sometimes, never.

Note that the order of these questions forces the respondents to evaluate their information level and get aware of their (lack of) information first. The questions on their confidence level are answered, thus, in full awareness of their available information. Combined the information and confidence questions indicate the level of objective ignorance (knowing/being aware information is lacking and still have confidence in your judgment).

The core of the survey consists of three cases testing the conjunction fallacy. Two of these cases consist of a case description followed by a question asking for a likelihood judgment (Cases 1 and 2). The third case is a replication of the original problem statement as used by Kahneman and Tversky. The context is reformulated to fit the security domain. As this reformulated problem shows the conjunction fallacy in plain sight, logic reasoning or recognition of the fallacy might influence the assessment of the respondents in the other two cases. Therefore, the reformulated problem is presented to the respondents as the third and final case study.

The reformulated problem consists of a short case description followed by a choice between two options. The respondents are asked to indicate which option they consider more likely. The first option has a general and short formulation. The second one is identical to the first option but is extended with more detailed information. Showing the two answers at the same time, in other words showing the conjunction rule, should or could guide the respondents to choose the shorter, more general, option. The second, more detailed, option, obviously is a sub-set of the first and should, therefore, be considered less likely.

The reformulated problem is kept identical for both the physical and cybersecurity community:

Case introduction:

Your organization is a large, international, pharmaceutical corporation based in the EU. Your R&D department has focused the last months on research in developing a COVID-19 vaccine. This department made considerable progress and is considered to be one of the global front runners and ahead of other
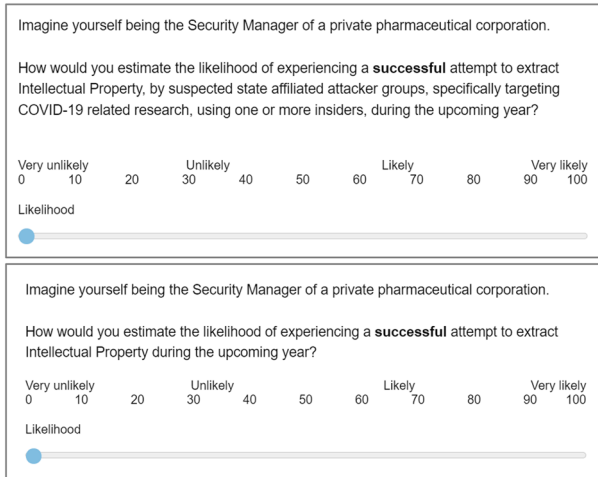
Fig. 2 Examples of the extended (top) and short (bottom) version of the security case experiment for the physical security domain showing the slider with the double scale

research institutes. Last week you discovered a serious attempt to steal information.

What is more likely:

- *This attack is launched by an organized crime organization.*
- *This attack is launched by an organized crime organization targeting IP (Intellectual Property) related to COVID-19 research.*

Note: this case is developed and presented to the respondents before in the real world COVID-19 vaccines were available. At the time the surveys were conducted in both the physical and cybersecurity domain several pharmaceutical corporations around the world were in the race of developing vaccines and there were indications (in the press) of attempts of IP theft at these kind of corporations. This case description can, therefore, be considered realistic.

Cases 1 and 2 are based on the same approach as the reformulated problem; however, in these two cases, the respondents are asked to estimate the likelihood of the case. Of each case there are two versions, a short and an extended version where three additional information elements are added. The respondents are automatically and randomly assigned to either the short or the extended version in a way that each respondent is offered one short version of an case and an extended version of the other. About half of the respondents first assessed the short version of Case 1 followed by the extended version of case 2 (group A). The other part of the respondents first assessed the extended version of case 1 followed by the short version of case 2 (group B). The likelihood estimation can be answered via a slider on a scale which offers the respondents both a probability scale (0–100%) and a qualitative likelihood scale (very unlikely, unlikely, likely, very likely). In Fig. 2 the short

and extended version of the same case are shown including the slider scale. After each case the respondents are asked to rate the importance of each information element for their likelihood assessment using a three point Likert scale (very important, important, not important). These two cases do not show or refer to the conjunction fallacy in any way. The respondents have no indication that they are offered a short or extended version.

To fit the two cases to the two domains, physical and cybersecurity, the description is adjusted to reflect domain specific realistic and recognizable cases. Case 1 is almost identical to the already discussed reformulated Problem (Case 3). The description of case 2 is made more specific for each domain. All case descriptions are based on real-life incidents or threats that were available in public sources (often in the press) at the time of conducting the surveys. Thus, they can be considered realistic. The structure of the cases and the number of additional detailed information aspects is identical for both domains. Table 1 shows all the case descriptions.

Finally the respondents are asked to express their expertise in a number of questions about individual characteristics. They are asked to indicate their age, number of years professional experience and number of years security experience. The current function of the respondents is asked including the number of years in this position. Finally they are asked to indicate their general education level (associate degree, bachelor degree or Master degree/PhD) and if any specific security trainings are completed. These individual characteristics may influence the individual assessments of the respondents.

The explorative results are retrieved via this online survey conducted between September 2020 and February 2021. Participation in the survey is promoted in both the IT and physical security professional community. It is promoted via LinkedIn and Twitter, both in general and in special interest groups like Security management, ASIS Europe and ASIS International, Dutch cybersecurity platform. Second, a direct email campaign is launched targeting the existing professional network of the researchers. Third, the survey is promoted via the Information Security Forum world conference: Digital 2020 (cybersecurity domain) and ASIS Europe 2021 conference (physical security domain). The sample of respondents ($N = 166$) is regarded a convenience sample.

Results and analysis

The results on the information position of the professionals are presented in Table 2. The security professionals indicate that, on average, about half the time they know the likelihood and consequences exactly. The respondents also indicate that they, on average, only sometimes, cannot estimate likelihood and consequences. One in four even indicates that they can always estimate likelihood and consequences, based on their experience and knowledge, even when they indicate they know they do not have accurate information.

Overall they claim to be confident about their judgment of likelihood and consequences most of the time (see Table 3).

Table 1 Case descriptions differentiated for the physical and cybersecurity domain, divided in group A and B (italic text indicates the three additional detailed information aspects in the extended case description)

Cybersecurity domain		Physical security domain	
Group A	Group B	Group A	Group B
Case 1 short	Case 1 extended	Case 1 short	Case 1 extended
Imagine yourself being the CISO of a private pharmaceutical corporation	Imagine yourself being the CISO of a private pharmaceutical corporation	Imagine yourself being the Security Manager of a private pharmaceutical corporation	Imagine yourself being the Security Manager of a private pharmaceutical corporation
How would you estimate the likelihood of experiencing a successful attempt to extract Intellectual Property during the upcoming year?	How would you estimate the likelihood of experiencing a successful attempt to extract Intellectual Property, *by suspected Chinese attacker groups, specifically targeting COVID-19 related research, using spear phishing techniques,* during the upcoming year?	How would you estimate the likelihood of experiencing a successful attempt to extract Intellectual Property during the upcoming year?	How would you estimate the likelihood of experiencing a successful attempt to extract Intellectual Property, *by suspected state affiliated attacker groups, specifically targeting COVID-19 related research, using one or more insiders,* during the upcoming year?
Case 2 extended	Case 2 short	Case 2 extended	Case 2 short
Imagine yourself being the CISO of a Fortune 500 corporation	Imagine yourself being the CISO of a Fortune 500 corporation	Imagine yourself being the Security Director of a fortune 500 logistics corporation	Imagine yourself being the Security Director of a fortune 500 logistics corporation
How would you estimate the likelihood of experiencing a successful attempt to execute a ransomware attack, *by criminal Russian hacker groups, using new targeted ransomware like Wasted-Locker, targeting the main ERP system (Enterprise Resource Planning),* during the upcoming year?	How would you estimate the likelihood of experiencing a successful attempt to execute a ransomware attack during the upcoming year?	How would you estimate the likelihood of experiencing a successful attempt *of bribery of employees of subcontractors, by organized crime organizations,* to facilitate international drug trafficking, *using maritime transport,* during the upcoming year?	How would you estimate the likelihood of experiencing a successful attempt to facilitate international drug trafficking, during the upcoming year?

Table 2 The information position of security professionals in security risk assessments

When evaluating security risks in general	Always (1)	Most of the time (2)	About half the time (3)	Sometimes (4)	Never (5)	Median answer	Mean answer*
I know the *likelihood* of security events *exactly*	2.0%	33.0%	19.3%	24.9%	20.8%	About half the time	3.29
I do not know the *likelihood* exactly but I have *quantified* information (evidence based probability)	4.6%	38.1%	23.4%	29.9%	4.1%	About half the time	2.91
I do not know the *likelihood* exactly but I *can estimate* the likelihood based on my experience and knowledge	9.6%	51.3%	23.9%	14.7%	0.5%	Most of the time	2.45
I do not know the *likelihood* exactly and I *cannot estimate* the likelihood based on my experience and knowledge	0.5%	13.7%	8.1%	54.3%	23.4%	Sometimes	3.86
I know the *consequences* of security events *exactly*	3.3%	42.4%	21.2%	19.6%	13.6%	About half the time	2.98
I do not know the *consequences* exactly but I have *quantified* information (evidence based probability)	3.3%	39.7%	20.7%	31.5%	4.9%	About half the time	2.95
I do not know the *consequences* exactly but I *can estimate* the likelihood based on my experience and knowledge	7.1%	49.5%	21.7%	19.6%	2.2%	Most of the time	2.60
I do not know the *consequences* exactly and I *cannot estimate* the likelihood based on my experience and knowledge	0.5%	12.0%	8.7%	50.5%	28.3%	Sometimes	3.94

*Considering the Likert scale a continues variable from always = 1 to never = 5

Table 3 Confidence levels of security professionals

When evaluating security risks in general	Always (1)	Most of the time (2)	About half the time (3)	Sometimes (4)	Never (5)	Median answer	Mean answer*
I feel confident about my assessments of the *likelihood* of security risks	8.3%	59.4%	20.0%	10.6%	1.7%	Most of the time	2.38
I feel confident about my assessments of the *consequences* of security risks	9.4%	64.4%	15.6%	9.4%	1.1%	Most of the time	2.28
I would feel *more confident* if I had more information on security risks	28.9%	33.3%	8.9%	27.8%	1.1%	Most of the time	2.39

*Considering the Likert scale a continues variable from always = 1 to never = 5

Individual characteristics might influence confidence. The respondents are asked to indicate their age, number of years professional experience, number of years security experience, the number of years in their current position, their general education level (associate degree, bachelor degree or Master degree/PhD) and if any specific security trainings are completed. To reduce this number of characteristics and explore their structure and influence all six items were subjected to an exploratory factor analysis with oblique rotation. The Kaiser–Meyer–Olkin measure verified the sampling adequacy for the analysis, KMO=0.741, Bartlett's test of sphericity v2 (15)=302.18, $p<0.005$, indicating that correlation structure is adequate for factor analyses.

Factor 1, reflecting experience, is comprised of four characteristics (age, number of years professional experience, number of years security experience, the number of years in their current position) that explain 44.4% of the common variance from all variables with factor loadings of 0.647 to 0.894. Factor 2 reflects specific security trainings and is comprised of one characteristic explaining 17.2% of the variance with a factor loading of 0.840. The final factor, reflecting education level, explains 16.6% of the variance with a factor loading of 0.840. All three factors have Kaiser's criterion of eigenvalues equal or greater than 1 and are sufficiently orthogonal to each other.

To assess the relationship between these factors and the confidence level of the respondents a Spearman's rank correlation is computed between the three factors and the three questions of Table 3.

Factor 1, experience, shows a negative correlation with the likelihood confidence level, $r(164)=-0.158$, $p=0.043$. This factor also shows a negative correlation with the consequence confidence level, $r(164)=-185$, $p=0.017$. Finally this factor shows a positive correlation with the confidence vs need for information level, $r(164)=0.229$, $p=0.003$. These results show that more experience significantly raises the number of occasions in which the respondents have confidence in their own assessments of likelihood and consequences. More experience, on the other hand, significantly reduces the number of occasions in which the respondents would require more information to be more confident.

No significant correlations are discovered between security specific trainings, factor 2, and confidence levels. These results indicate that completing security specific trainings do not influence the level of confidence of the respondents in their own assessments.

The third and final factor, education level shows a significant positive correlation with the likelihood confidence level, $r(164)=0.179$, $p=0.021$, and the consequence confidence level, $r(164)=0.239$, $p=0.002$. No significant correlation is noted between factor 3 and the confidence vs need for information level. A higher education level, thus, leads the respondents to less occasions in which they are confident about their assessments of likelihood and consequences.

Table 4 shows the combined results of the first knowledge question as it asked for the most exact information (Table 2) and the confidence questions. A normative assumption might be that respondents that indicate to have exact information can be expected to be confident about their assessments and the opposite. Following this assumption the diagonal from the upper left corner (always exact knowledge

Table 4 Information vs confidence levels of security professionals (in number of respondents)

When evaluating security risks in general:	I feel confident about my assessments of the **likelihood** of security risks					
Note: number of respondents	Always	Most of the time	About half the time	Some-times	Never	**Total:**
I know the **likelihood** of security events exactly:						
Always	2	1	-	-	-	3
Most of the time	11	45	7	1	-	64
About half the time	2	22	8	1	0	33
Sometimes	-	25	13	2	1	41
Never	-	14	8	15	2	39
Total:	15	107	36	19	3	180

When evaluating security risks in general:	I feel confident about my assessments of the **consequences** of security risks					
Note: number of respondents	Always	Most of the time	About half the time	Some-times	Never	**Total:**
I know the **consequences** of security events exactly:						
Always	2	4	-	-	-	6
Most of the time	14	56	6	1	-	77
About half the time	1	26	11	1	-	39
Sometimes	-	22	6	4	1	33
Never	-	8	5	11	1	25
Total:	17	116	28	17	2	180

and always confident) to the lower right corner (never exact knowledge and never confident) show the respondents which seem to align their knowledge and confidence. As stated in the introduction exact knowledge on future events is considered intractable knowledge. The respondents in the dotted oval, more than half of the respondents (likelihood: 54.4%, consequences: 67.2%), thus, seem to overestimate their knowledge. The lower left area (gray) contains respondents confirming to lack exact information most often but are often confident about their assessments. These respondents (likelihood: 33.3%, consequences: 22.8%) seem to show objective ignorance being more confident than their information position would permit.

Case 1

Figure 3 shows the results of case 1 (the results of both the physical and cybersecurity domains are combined). Professionals working in the same domain with comparable general knowledge reach, based on identical information, likelihood assessments ranging from 0 to 100% for the short version of case 1 ($n = 90$) and 10 to 100% for the extended version of case 1 ($n = 87$).

The median answer for case 1 short is 65%, the average answer is 57.1% ($M = 57.1$, $SD = 26.33$, $Q1 = 32.5\%$, $Q3 = 80\%$). The median answer for case 1 extended is 75%, the average answer is 69.6% ($M = 69.6$, $SD = 21.56$, $Q1 = 60\%$,

Fig. 3 Results of likelihood assessments of case 1

Q3 = 84%). An independent sample *T*-test is conducted to compare these assessments: the identified average difference of 12.5% is significant, $t(175) = -3.449$, $p = 0.001$.

The group of respondents assessing the extended case, including specific conditions, estimated the likelihood on average at 69.6% while the group assessing the short version of the same case, thus, without specific conditions, estimated the likelihood 57.1%. This significant mean difference seems to express the effect of the conjunction fallacy (the assumption that more specific conditions are more probable).

Case 2

Figure 4 shows the results of case 2. The results of this case show almost no influence of the conjunction fallacy. The average likelihood assessment of the case 2 extended option is only slightly higher ($M = 57.5\%$, SD = 24.43, $n = 87$) than the average likelihood assessment of the case 2 short option ($M = 56.3\%$, SD = 23.83, $n = 84$). This difference is not significant.

The results of case 1 seem to show the effects of the conjunction fallacy while the results of case 2 do not. This different average reaction to these two cases can be caused by either the difference between the content of the cases (the structure and number of specific conditions of the two cases is identical) and/or a possible difference between the two randomly assigned groups. The difference in content of the two cases will be analyzed in the discussion section. As the structure of the two cases is identical for this section we assume they would evoke comparable reactions.
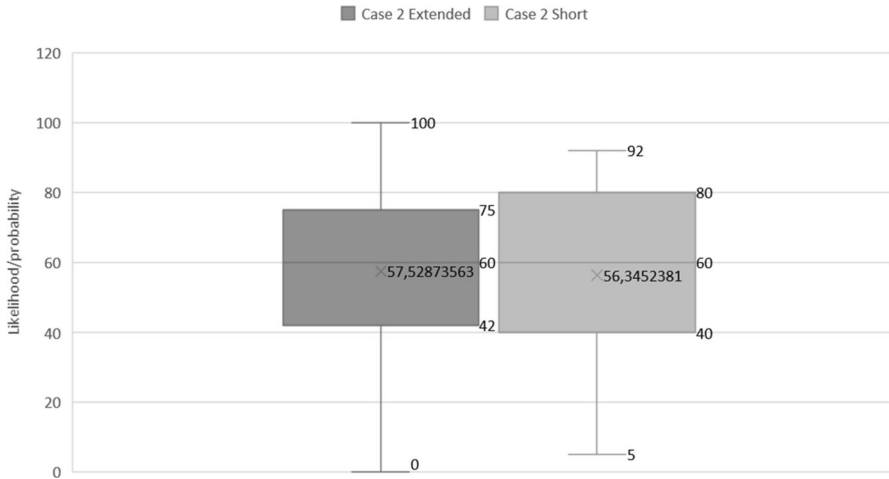
Fig. 4 Results of likelihood assessments of case 2

Table 5 Comparing characteristics of randomly composed groups A and B

		Group A	Group B
Average likelihood assessment (case 1 & 2)	Short case description	57.13% (1)	56.34% (2)
	Extended case description	57.52% (2)	69.63% (1)
Combined average likelihood assessment		57.33%	63.11%
N		85	81
Age (in years)		49.1	50.2
Total professional experience (in years)		25.7	24.4
Total security experience (in years)		18.3	17.9
Current position (in years)		8.5	7.6
Education level	Associate degree	17.6%	13.6%
	Bachelor degree	40.0%	43.2%
	Master degree/PhD	42.4%	43.2%
Security specific training		62.4%	70.4%
Case 3 'reformulated problem'	Short answer	25.0%	25.9%
	Extended answer	75.0%	74.1%

Table 5 shows the composition of the two groups in which the group of respondents confronted with the short version of case 1 first followed by the extended version of case 2 is denoted as group A. Group B assessed the extended version of case 1 first followed by the short version of case 2.

Both groups reacted similar to the conjunction fallacy as presented in the reformulated problem (case 3). Three out of four respondents of both groups selected the answer with more specific conditions and thus show vulnerability for the conjunction fallacy.

There are also no significant differences between average individual characteristics of the two groups. As there seems to be no indication for a difference between the groups and they are equally vulnerable to the conjunction fallacy, we might expect comparable risk assessments.

Combining the average likelihood assessments of the two cases for each individual respondent shows an average for the respondents in group A of 57.33% while the respondents in group B on average assess the likelihood 63.11%. On average group B estimates the likelihood of the combined two cases 5.8% higher (absolute difference) which is a relative difference of 9.2%.

Case 3 security conjunction: the reformulated problem

A total of 165 respondents answered the reformulated problem. 42 (25.5%) considered the first (short) option more likely, 123 (74.5%) the second (extended) one. In the physical security domain 58.8% of the respondents followed the fallacy and choose the extended option. Of the respondents active in the cybersecurity domain even 81.6% selected the extended option.

Discussion and conclusions

On average the respondents indicate that they have exact or quantified information about likelihood and consequences about half the time. This finding deviates from the expectation that security professionals would recognize their information position about security risks as both imperfect and intractable. However, they also indicate that they can estimate the likelihood and consequences most of the time (and only sometimes cannot estimate at all). Assuming that the respondents are right about their knowledge position they assess risk half of the time based on information (evidence based). On the other hand they assess security risks without proper information also half of the time and still come up with an estimation of likelihood and consequences. As these assessments have a serious impact on security risk decision making and the allocation of resources to manage, mitigate, and/or accept these risks, it is worth noting that these decisions do not seem to be based on evidence about half of the time.

The perception of the respondents on their information position can be questioned. As risk assessments are in fact predictive judgments and the information about the future can be considered intractable by nature, this perception of the security professionals can be considered audacious.

Overall the majority of the security professionals in this study indicate that they are always or most of the time confident about their assessments (for likelihood assessments 67.7%, for consequence assessments 73.8%). This level of confidence can be considered in agreement with the information position considering the perceived information position of the professionals as indicated above. It was hypothesized that the security professionals would show modest confidence based on the assumption that exact and/or evidence based information on security risks is often

lacking. They, however, seem to ignore the latter and thus show a higher level of confidence than expected. As the respondents on average indicate to hold exact or quantified information only half of the time, they, thus, might be considered over-confident about their risk assessments. Combining the perceived information position of the professionals with their confidence reveals objective ignorance. A portion of respondents indicate they have exact information only sometimes or even never but are confident most or half of the time (for likelihood assessments 33.3%, for consequence assessments 23.3%). These respondents are aware of their lack of exact information but are confident nevertheless. This lack of information does not seem to affect their ability to form a predictive judgment and be confident about it.

Individual characteristics influence confidence levels. As hypothesized more professional and security experience significantly raises the confidence level of the security professionals. More experienced security professionals are more often confident about their assessments of both likelihood and consequences. More experienced security professionals also indicate that more information would raise their confidence level to a lesser extent than less experienced professionals indicate. In short these results seem to indicate that more experience leads to higher levels of (over)confidence and less need for additional information. These findings confirm results previous work (Desender et al. 2018; Sieck and Yates 1997). A higher education level on the other hand significantly reduces the confidence in likelihood and consequences assessments. These results might prove the adage 'the more you know, the more you realize you don't know' as other scholars also found (Wright and Ayton 1986). Security specific trainings do not significantly influence confidence level or the need for additional information.

The third case (reformulated problem) in this study clearly proved the significant influence of more detailed information on likelihood assessments as expected. Three in four of the security professionals assess the likelihood of a more detailed case higher. This case offered the two answer options in one single view, showing the conjunction fallacy in plain sight. This, however, did not lead the majority of the professionals to apply logical reasoning and select the option with the shorter description. These results replicate numerous previous studies in other domains showing the power of details, stories, and assumptions. This study, for the first time, shows this effect on a realistic real-life security risk case.

The significant effects of the conjunction fallacy on security risk likelihood assessments are visible in the results of case 1. The likelihood of the short case is on average estimated at 57.1% while the likelihood of the extended version is estimated at 69.6%. In contrast to the expectation it is worth to note that the assessments of the security professionals, with similar backgrounds, professions, and experience, show a substantial variance or so called system noise (short case description: $M = 57.1\%$, SD $= 26.33\%$, extended case description: $M = 69.6\%$, SD $= 21.56\%$). Even with the presented limited case descriptions their assessments of the likelihood vary from unlikely to very likely. As these security professionals each decide or influence security risk decision making in their own organization, these results denote the possible variation in response to similar risks between different organizations.

The likelihood assessments of the two groups at case 2 show different results compared to case 1. There is hardly any difference in the likelihood assessment

of the short case description ($M=56.3\%$, SD$=23.83\%$) and the assessment of the extended case description ($M=57.5\%$, SD$=24.43\%$). The level of system noise is similar to case 1.

As the two randomly assigned groups do not significantly differ in characteristics (see Table 5), the difference between the likelihood assessments of cases 1 and 2 can only be caused by either the experiment setup and/or the different subject/content of the cases. In the following several possible explanations for the difference in overall response form group A and B are discussed.

The characteristics of the respondents in the two groups do not differ significantly; however, their average assessment of the two cases combined shows a significant difference. The average assessment of the two cases is 5.8% point higher in group B compared to group A. One of the possible explanations for this difference could be so called level noise, variability of judgment between individuals (fe. some security professionals might be more risk averse than others). Correcting the average assessments of the two cases for this possible level noise would lead to an average difference between the short and extended versions at case 1 of 6.7% point and for case 2 of 7% point. In both cases the extended version is assessed a comparable higher likelihood. Assuming this reasoning valid the conjunction fallacy raises the likelihood assessment with 6.7–7% point.

The setup of the experiment led the respondents to first assess case one followed by case two. As a consequence group A was first presented a short description of case 1 followed by an extended description of case 2. Group B, on the other hand, was confronted with first an extended description (case 1) followed by a short case description (case 2). The assessments of the first case might influence the respondents at their assessment of the second case, for example by the anchoring effect. This cognitive bias points at a human tendency to focus on a first piece of information to make subsequent judgments. Even if this piece of information is not related to the following judgment, this 'anchor' is proven to be influential. In this case the first assessment might become an anchor for the second assessment. We observe almost no difference in the average likelihood assessments over all group A respondents for the short and extended case study descriptions (57.1% vs 57.5% resp), which might suspect an anchoring effect, although no definitive proof can be given for such effect based on the current data. The average likelihood assessments over all group B respondents for the short and extended case study descriptions does show a large difference (69.6% vs. 56.3% resp), but also here no definitive proof can be given that there is absence of the anchoring effect. There might be other factors which influence the difference in the average likelihood assessments over the group respondents for the short and extended case study descriptions.

The two cases each describe a realistic, actual, real-life security risk. The first case describes a situation which, at the time of the experiments, was very relevant and discussed publicly. The second case is as relevant and actual as the first but was less prominent. The difference between the results of the two domains might be explained by the theory of hints (Kohlas and Monney 2013). Previous work by Brachinger and Monney explains the fallacious behavior of individuals as indicated by the conjunction fallacy (Brachinger and Monney 2003). In their study they show that individuals confronted with a choice, in which only vacuous mindless hints and no

precise hints are available, are forced to refer to their general knowledge to retrieve a subjective probability. In such situations the subjective interpretation of simple hints guides the decision maker. In this study both case introductions contain only vacuous hints. None of these hints indicates any precise information about the likelihood of interest by an organized crime organization, the target Intellectual Property (IP) or even more specific IP related to COVID-19 research. The simple (supporting) hints in the introduction about the position on the development of a COVID-19 vaccine at the hypothetical pharmaceutical corporation, might imply a large value at stake leading to interest of various malicious actors like organized crime. These simple hints can also lead to the interpretation that the most obvious information to extract is IP related to COVID-19 research. Other possible, and equally realistic, options like an attempt to extract commercial information by a foreign competitor or state affiliated actor might be discarded by the respondents. The same arguments apply on the second case of which the structure is similar.

Forcing the respondents to refer to their individual frame of reference, prior experience or expertise, as this theory stresses, can explain the difference between the results in between the two cases. The first case related to very prominent and available information and discussion while for the subject of the second case was less attention at that point in time.

This theory might also explain the difference in response between the physical en cybersecurity domain at case 3. In the physical security domain 58.8% of the respondents followed the fallacy and chose the extended option. Of the respondents active in the cybersecurity domain 81.6% selected the extended option. Both the domains are closely related but deal with different threats. As an indication: the top threat in the cybersecurity domain in 2020 was IP theft by various threat vectors (ISACA 2020) while in the physical security domain the top threat in 2020 was malicious physical access (ENISA 2020). The respondents originating from the cybersecurity domain, therefore, might relate more to option: 'organized crime organization targeting IP related to COVID-19 research.' It fits their frame of reference, might lead to a stronger representativeness, recognition and emotion, and thus, availability. According to the theory of hints and the study of Brachinger and Monney this explains the fall for the conjunction fallacy. An important consequence of this conclusion can be that professionals with domain expertise, and thus a deeper subjective interpretation of simple hints, and readily available information or even experience (Dumm et al. 2020), assess a higher likelihood to risks in their domain than non-domain experts.

In agreement with the hypothesis the results of this study clearly show the influence of the conjunction fallacy on the judgment of security professionals. The consequence of this fallacy in the security domain can influence security risk assessments by these practitioners considerably. Following the fallacy, retrieving more specific, detailed and recognizable information may lead the individual professional to consider a case, incident, or threat more likely which in turn might lead to distorted risk assessments in organizations and society. Security professionals, facing the difficult daily task to assess security risks, often based on little accurate information, seem to be confident about their predictive judgment. This study hopes to raise awareness for possible flaws, unknown overconfidence, and ignorance of security

professionals. As a whole, these findings have important implications for the professional security community and anyone depending on it.

Declarations

Conflict of interest The authors report no conflict of interest.

References

Allingham, M. (2002). *Choice theory: A very short introduction*. OUP Oxford.

Alruwaii, A., and D.J. Brooks. 2008. Organisational security: A propositional study to map expert knowledge. Paper presented at the proceedings of the 1st Australian Security and Intelligence Conference.

Andersson, O., H.J. Holm, J.-R. Tyran, and E. Wengström. 2020. Robust inference in risk elicitation tasks. *Journal of Risk and Uncertainty* 61 (3): 195–209.

ANSI/ASIS. 2012. *Security management standard: Physical asset protection*. Alexandria: ASIS International.

ANSI/ASIS/RIMS. 2015. *Risk assessment RA1.2015*. Alexandria: ASIS International.

ASIS_International. 2015. *Risk assessment, ANSI/ASIS/RIMS RA.1-2015*. Alexandria: ASIS International.

Baron, J. 2004. *Normative models of judgment and decision making*. Hoboken: Wiley Online Library.

Bonini, N., K. Tentori, and D. Osherson. 2004. A different conjunction fallacy. *Mind & Language* 19 (2): 199–210.

Brachinger, H.W., and P.A. Monney. 2003. The conjunction fallacy: explanations of the linda problem by the theory of hints. *International Journal of Intelligent Systems* 18 (1): 75–91.

Carbone, E., X. Dong, and J. Hey. 2017. Elicitation of preferences under ambiguity. *Journal of Risk and Uncertainty* 54 (2): 87–102.

Charness, G., T. Garcia, T. Offerman, and M.C. Villeval. 2020. Do measures of risk attitude in the laboratory predict behavior under risk in and outside of the laboratory? *Journal of Risk and Uncertainty* 60 (2): 99–123.

Cooke, R.M. 1991. *Experts in uncertainty*. New York: Oxford University Press.

de Wit, J., W. Pieters, S. Jansen, and P. van Gelder. 2021. Biases in security risk management: Do security professionals follow prospect theory in their decisions? *Journal of Integrated Security and Safety Science* 1 (1): 34–57.

Desender, K., A. Boldt, and N. Yeung. 2018. Subjective confidence predicts information seeking in decision making. *Psychological Science* 29 (5): 761–778.

Dumm, R.E., D.L. Eckles, C. Nyce, and J. Volkman-Wise. 2020. The representative heuristic and catastrophe-related risk behaviors. *Journal of Risk and Uncertainty* 60 (2): 157–185.

ENISA. 2020. Physical manipulation, damage, theft, loss. *ENISA Threat Landscape*. https://www.enisa.europa.eu/publications/physical-manipulation-damage-theft-loss

Fantino, E., J. Kulik, S. Stolarz-Fantino, and W. Wright. 1997. The conjunction fallacy: A test of averaging hypotheses. *Psychonomic Bulletin & Review* 4 (1): 96–101.

Fiedler, K. 1988. The dependence of the conjunction fallacy on subtle linguistic factors. *Psychological Research Psychologische Forschung* 50 (2): 123–129.

Gigerenzer, G. 1991. How to make cognitive illusions disappear: Beyond "heuristics and biases." *European Review of Social Psychology* 2 (1): 83–115.

Gigerenzer, G., and R. Selten. 2002. *Bounded rationality: The adaptive toolbox*. Cambridge: MIT Press.

Hansson, S.O. 2012. A panorama of the philosophy of risk. In *Handbook of risk theory: Epistemology, Decision Theory, Ethics, and Social Implications of Risk, 1*, 27–54. Dordrecht: Springer Science+Business Media B.V.

Hertwig, R., and G. Gigerenzer. 1999. The 'conjunction fallacy' revisited: How intelligent inferences look like reasoning errors. *Journal of Behavioral Decision Making* 12 (4): 275–305.

Information_Security_Forum. 2018. *Standard of good practice*. Surrey: Information Security Forum.

ISACA. 2020. *Top Cyberattacks of 2020 and How to Build Cyberresiliency.* https://www.isaca.org/resources/news-and-trends/industry-news/2020/top-cyberattacks-of-2020-and-how-to-build-cyberresiliency

ISO. 2018. *ISO 31000 risk management—Guidelines.* Geneva: International Organization for Standardization.

ISO. 2019. *ISO 22301 security and resilience—Business continuity management systems—Requirements.* Geneva: International Organization for Standardization.

ISO/IEC. 2011. *ISO/IEC 27005 Information technology_Security_techniques_Information security risk management.* Geneva: ISO.

Jain, K., K. Mukherjee, J.N. Bearden, and A. Gaba. 2013. Unpacking the future: A nudge toward wider subjective confidence intervals. *Management Science* 59 (9): 1970–1987.

Kahneman, D. 2012. *Ons feilbare denken: Thinking, fast and slow.* Business Contact.

Kahneman, D., O. Sibony, and C.R. Sunstein. 2021. *Noise, a flaw in human judgment.* London: William Collins.

Kohlas, J., and P.A. Monney. 2013. *A mathematical theory of hints: An approach to the Dempster-Shafer theory of evidence* (Vol. 425): Dordrecht: Springer Science+Business Media B.V.

Kuhn, K.M., and J.A. Sniezek. 1996. Confidence and uncertainty in judgmental forecasting: Differential effects of scenario presentation. *Journal of Behavioral Decision Making* 9 (4): 231–247.

Ludwin-Peery, E., N.R. Bramley, E. Davis, and T.M. Gureckis. 2020. Broken physics: A conjunction-fallacy effect in intuitive physical reasoning. *Psychological Science* 31 (12): 1602–1611.

Möller, N. 2012. The concepts of risk and safety. In *Handbook of risk theory: Epistemology, decision theory, ethics, and social implications of risk*, 55–85. Dordrecht: Springer Science+Business Media B.V.

Riesch, H. 2013. Levels of uncertainty. In *Essentials of risk theory*, ed. S. Roeser, R. Hillerbrand, P. Sandin, and M. Peterson, 29–56. Dordrecht: Springer.

Sieck, W., and J.F. Yates. 1997. Exposition effects on decision making: Choice and confidence in choice. *Organizational Behavior and Human Decision Processes* 70 (3): 207–219.

Simon, H.A. 1982. *Models of bounded rationality: Empirically grounded economic reason*, vol. 3. Cambridge: MIT Press.

Slovic, P.E. 2000. *The perception of risk.* London: Earthscan Publications.

Slovic, P. 2010. The feeling of risk. In *New perspectives on risk perception*. New York: Routledge

Stolarz-Fantino, S., E. Fantino, D.J. Zizzo, and J. Wen. 2003. The conjunction effect: New evidence for robustness. *American Journal of Psychology* 116 (1): 15–34.

Sunstein, C.R. 2005. *Laws of fear: Beyond the precautionary principle.* Cambridge University Press.

Talbot, J., and M. Jakeman. 2011. *Security risk management body of knowledge*, vol. 69. Hoboken: Wiley.

Tentori, K., N. Bonini, and D. Osherson. 2004. The conjunction fallacy: A misunderstanding about conjunction? *Cognitive Science* 28 (3): 467–477.

Tentori, K., and V. Crupi. 2012. On the conjunction fallacy and the meaning of and yet again: A reply to. *Cognition* 122 (2): 123–134.

Tversky, A., and D. Kahneman. 1975. Judgment under uncertainty: Heuristics and biases. In *Utility, probability, and human decision making*, ed. D. Wendt and C. Vlek, 141–162. Dordrecht: Springer.

Tversky, A., and D. Kahneman. 1983. Extensional versus intuitive reasoning: The conjunction fallacy in probability judgment. *Psychological Review* 90 (4): 293.

Tversky, A., and D. Kahneman. 2008. Extensional versus intuitive reasoning: The conjunction fallacy in probability judgment. *Reasoning: Studies of human inference and its foundations,* 114-135.

Tversky, A., and D.J. Koehler. 1994. Support theory: A nonextensional representation of subjective probability. *Psychological Review* 101 (4): 547.

Wright, G., and P. Ayton. 1986. Subjective confidence in forecasts: A response to Fischhoff and MacGregor. *Journal of Forecasting* 5 (2): 117–123.