Shreeve, B., Hallett, J., Edwards, M., Anthonyamy, P., Frey, S., & Rashid, A. (2020). "So if Mr Blue Head here clicks the link...": Risk thinking in cyber security decision making. *ACM Transactions on Privacy and Security*, 24(1), [5]. <https://doi.org/10.1145/3419101>

Peer reviewed version

Link to published version (if available):
[10.1145/3419101](https://doi.org/10.1145/3419101)

[Link to publication record in Explore Bristol Research](#)
PDF-document

This is the author accepted manuscript (AAM). The final published version (version of record) is available online via ACM at <https://doi.org/10.1145/3419101>. Please refer to any applicable terms of use of the publisher.

University of Bristol - Explore Bristol Research

General rights

This document is made available in accordance with publisher policies. Please cite only the published version using the reference above. Full terms of use are available:
<http://www.bristol.ac.uk/red/research-policy/pure/user-guides/ebr-terms/>

“So if Mr Blue Head here clicks the link...” Risk thinking in cyber security decision making

BENJAMIN SHREEVE, University of Bristol

JOSEPH HALLETT, University of Bristol

MATTHEW EDWARDS, University of Bristol

PAULINE ANTHONYSAMY, Google

SYLVAIN FREY, Google

AWAIS RASHID, University of Bristol

Cyber security decision making is inherently complicated, with nearly every decision having knock-on consequences for an organisation’s vulnerability and exposure. This is further compounded by the fact that decision-making actors are rarely security experts, and may have an incomplete understanding of the security that the organisation currently has in place. They must contend with a multitude of possible security options that they may only partially understand. This challenge is met by decision-makers’ *risk thinking*—their strategies for identifying risks, assessing their severity, and prioritising responses. We study the risk thinking strategies employed by teams of participants in an existing data set derived from a tabletop cyber-physical systems security game [16]. Our analysis identifies four *structural patterns* of risk thinking and two *reasoning* strategies: *risk-first* and *opportunity-first*. Our work highlights that risk-first approaches (as prescribed by the likes of NIST-800-53 [22] and ISO27001 [21]) are followed neither substantially nor exclusively when it comes to decision-making. Instead, our analysis finds that decision-making is affected by the plasticity of teams: that is, the ability to readily switch between ideas and practising both risk-first and opportunity-first reasoning.

CCS Concepts: • **Social and professional topics** → **Computing occupations; User characteristics; • Applied computing** → **Decision analysis.**

Additional Key Words and Phrases: Decision-Making, Cybersecurity, Cybersecurity Professions

ACM Reference Format:

Benjamin Shreeve, Joseph Hallett, Matthew Edwards, Pauline Anthonysamy, Sylvain Frey, and Awais Rashid. 2020. “So if Mr Blue Head here clicks the link...” Risk thinking in cyber security decision making. *ACM Trans. Priv. Sec.* 1, 1, Article 1 (January 2020), 29 pages. <https://doi.org/10.1145/3419101>

1 INTRODUCTION

Professional frameworks, such as ISO 27001 [21], highlight risk as a key component of cyber security decision making. Such approaches assume that individuals or teams actually practice risk thinking or that such risk thinking is uniform—often rooted in definitions such as Risk = Threat × Vulnerability × Consequences. However, little is known about the types of risk thinking utilised by different demographics (e.g., managers, IT personnel, security experts) and their reasoning when making cyber security decisions. Our work in this paper, is motivated by a need to better

Authors’ addresses: Benjamin Shreeve, ben.shreeve@bristol.ac.uk, University of Bristol; Joseph Hallett, joseph.hallett@bristol.ac.uk, University of Bristol; Matthew Edwards, matthew.john.edwards@bristol.ac.uk, University of Bristol; Pauline Anthonysamy, pauline.anthonysamy@gmail.com, Google; Sylvain Frey, frey.sylvain@gmail.com, Google; Awais Rashid, awais.rashid@bristol.ac.uk, University of Bristol.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

© 2020 Copyright held by the owner/author(s).

2471-2566/2020/1-ART1

<https://doi.org/10.1145/3419101>

understand how teams go about making cyber security decisions and the role risk thinking plays in this.

Research from risk theory suggests that people utilise risk thinking in order to avoid ambiguity. Ambiguity is sometimes categorised in terms of knowability [8]; that is, the extent of one's confidence in knowledge. Wang and Nyshadham [48] have adapted this to explore how e-commerce consumers evaluate risk—they categorise a person's knowledge of a risk as falling under one of four states:

- *known certainty*: Where the decision maker has information on all aspects of the decision—there is nothing that is unknown or unknowable. For example, an organisation knows who is going to attack it, how it is going to be attacked, what is its current exposure to the attack and multiple options are available to defend against the attacks.
- *known uncertainty*: Where there is a probability associated with the risk but it has a specified (and agreed upon) value. For example, an organisation has calculated that an attack would have a 10% chance of success with a known impact.
- *unknowable uncertainty*: Where the situation is accepted as so complex that no one can know everything and no risk probability can be effectively calculated or estimated. In this situation an organisation wouldn't be certain about who was going to attack it, or how, nor what is its current exposure.
- *unknown uncertainty*: A situation where the risk probability isn't known by all but where it may be known to some. For example, an organisation has calculated that an attack has a chance of success of less than 15% but the actual likelihood isn't known nor is the impact.

These four framings help to explain how decision makers can think about risk. The majority of cyber security decisions are often made under *unknowable uncertainty* or *unknown uncertainty*. That is, it is often impossible to account for all of the consequences of a cyber security decision.

In this work we set out to explore the risk thinking practices that people employ when faced with such uncertainty. We analyse an existing data set of cyber security decisions from Frey et al. [16]. The data set involves 18 hours of recordings, from 12 teams playing a tabletop cyber security game. Recordings are transcribed as 12,846 paragraphs of dialogue, of which 1,390 paragraphs are related to risk thinking. The game used (Decisions & Disruptions (D-D)) emulates some of the complexities of cyber security decision making by challenging teams to help a fictitious hydro-electric company develop cyber security defences. Teams are provided with a Lego representation of the company's existing infrastructure and assets, given a finite budget and a number of security controls and information gathering activities which they may purchase. They suffer a range of typical cyber attacks that may arise as a result of their choices.

Teams have to collaborate, sharing information and ideas in order to identify the security controls or information gathering options in which to invest. Their conversation during this process is recorded and transcribed. We then capture their *risk thinking* via heat maps and graph-like structures—identifying structural patterns as well as reasoning approaches that lead to the decisions the teams make when playing the game. Whereas Frey et al. [16] have studied *what* decisions were necessary to bring about a good security outcome in the context of the game, we study *why* teams arrive at their particular choices and *how* they do so, (i.e., the structures and reasoning that lead up to the decisions).

Our analysis offers a number of insights into risk thinking practices:

- Discussion of specific *assets* at risk was the least common, being less than 1% of all risk discussion. In contrast, teams considered the potential *vulnerabilities* of the organisation the most during their risk discussion, this being more than twice as common as the discussion of *impact* (the next most common risk discussion category) or *threats*. By raising awareness of this bias

we hope to improve decision-making diversity by encouraging decision-makers to increase the extent to which they consider assets, threats and impact, and not just vulnerabilities primarily.

- Discussion around risk was largely front-loaded, with the majority of risk related dialogue taking place during the first round of the game (as teams worked to understand the game or game scenarios). Use of risk thinking then suffered a gradual decline over the later rounds before coming back to the fore when teams entered the final round and became aware that they only had one ‘play’ remaining. Understanding the natural location of risk thinking during decision making can help in the formation of prompts to help decision-makers think about risk at the right moment during the decision-making process.
- Teams’ risk thinking structures were primarily characterised by simple forms of thinking—*isolated*, *sequential* or *radial*—with very little utilisation of *complex thinking*. *Isolated* thinking is where parts of the conversation may occur one after the other but where there is no obvious relation between them. *Sequential* thinking describes how one point in a conversation leads naturally onto another point. *Radial* thinking is where a point prompts the generation of multiple ideas (e.g., brainstorming). *Complex* thinking meanwhile describes conversation which cross-references ideas and reflects on past suggestions and decisions. This finding is particularly interesting given the extent to which the practitioner community relies on highly-structured methods for capturing risk-thinking. Further research is needed on this point to explore the impact that structured approaches have on decision-making vs less-structured approaches.
- Teams utilise two forms of risk-focused reasoning during their decision-making: *risk-first* whereby they identify risks and then seek to discover the optimum mitigation; and *opportunity-first*, whereby they first consider the investment or opportunity available and then evaluate its effectiveness in mitigating potential risks or giving rise to new ones. Teams do not appear to follow one reasoning approach exclusively—switching between them. The fact that teams left to their own devices (i.e., not told to use a specific threat model or approach) utilise multiple approaches brings into question why a risk-first approach is always the dominant decision-making approach promoted within the practitioner community.

The novel contribution of our work is in this exploration of whether and how teams practice risk thinking and the forms of reasoning approaches they use when encountering cyber security risks. There is an increasing focus on risk assessment and management in the context of cyber security. We understand little, however, of the kind of risk-thinking patterns and reasoning approaches that those charged with such risk assessments may use and whether these lead to an effective exploration of the range of risks to the organisation and its infrastructure. Ours is the first to undertake a detailed analysis of teams from different backgrounds making cyber security decisions—specifically analysing the patterns of risk thinking (or lack thereof) and the reasoning approaches (or lack thereof) exhibited by the teams. These findings can form the basis of guiding risk decision-makers to avoid isolated and sequential thinking when making cyber-risk decisions and consider a combination of reasoning approaches as reflected by opportunity-first and risk-first reasoning in our analysis. This can be achieved, for instance, through reconsideration and adaptation of the very design and steps used in risk analysis frameworks—so that they encourage radial and complex thinking as well as balance between opportunity-first and risk-first reasoning to ensure that the business goals and cyber security needs remain balanced.

Fig. 1. D-D starting layout

Table 1. Cost of each investment in the game. Items with a ★ are only available *after* the *Asset Audit* has been played.

Investment	Cost
Encryption (PCs)★	\$20,000
Encryption (databases)★	—
Threat Assessment	—
Antivirus	\$30,000
Asset Audit	—
Controller Upgrade★	—
Firewall (Office)	—
Firewall (Plant)	—
PC Upgrade★	—
Security Training	—
Server Upgrade★	—
CCTV (Office)	\$50,000
CCTV (Plant)	—
Network Monitoring (Office)	—
Network Monitoring (Plant)	—

2 METHOD

2.1 The game

We use D-D [16], a tabletop game (see Figure 1) where a team of 2-6 players are put in charge of the security at a small hydro-electric power company, under the direction of a Game Master who enforces the game's rules. The company is divided into two physical sites connected to the Internet via routers: the *power plant* itself, where water turbines are controlled by a SCADA controller, which sits on a LAN along with a database of production data and engineer PCs; and the company *office* that hosts a web and email server, business database, and employee PCs on its own LAN.

Teams play through 4 rounds, and are given a limited budget of \$100,000 per round (with any unused rolling between rounds) to spend on a range of potential cyber defences (see table 1). At the end of each round, after new defences are installed, the Game Master describes a number of attacks

Table 2. Group names and player distribution

	Academia	Industry
Security Experts	SA1 (4 PhD students) SA2 (3 undergrad. stud.)	SI1 (4 consultants) SI2 (5 consultants)
Computer scientists	CA1 (2 academics) CA2 (4 postgrad. stud.)	CI1 (6 IT engineers) CI2 (4 IT engineers)
Managers	MA1 (3 postgrad. stud.) MA2 (4 undergrad. stud.)	MI1 (2 managers) MI2 (2 managers)

happening, that may or may not be stopped depending on the defensive choices of the players. There are 33 predetermined attacks that teams may suffer, Frey et al. [16] identified this set of attacks through extensive play-testing with cyber security practitioners. Two of the potential investment options (*threat assessment* and *asset audit*) can provide the teams with immediate additional intelligence—to inform their decisions. The *asset audit* reveals a number of new vulnerabilities and results in additional investment options becoming available to address these.

Because money per round is limited, the players must prioritise their investments to address urgent threats first while delaying defences against less likely or less impacting threats to the later rounds of play. The players must achieve consensus on investment decisions, based on the Game Master’s input and their own intuition and experience. The discussions to get such consensus, which the Game Master encourages but does not influence outside of the inputs defined by the game rules, are, therefore, a direct dive into the players’ perception of risk and security.

2.2 Data Set

We utilise the data set originally gathered by Frey et al. [16], pertaining to 12 teams playing the D-D decision-making game (see table 2). The 12 teams represent different demographics: Teams *SA1*, *SA2*, *SI1* and *SI2* are cyber security experts—participants had skills, degrees and/or professional cyber security experience; Teams *CA1*, *CA2*, *CI1* and *CI2* are more generalist computer scientists—they had a background in computer science, but not in cyber security; and teams *MA1*, *MA2*, *MI1* and *MI2* are managers—they were all from a management background and had no skills in computer science nor cyber security. The ‘A’ and ‘T’ variations denote teams recruited from *Academia* and *Industry* respectively. We discuss the impact that these different backgrounds have on decision-making later on in Section 3.3.

Frey et al. [16] originally used the data set to evaluate the *quality* of the participants’ decisions—looking to measure how well the different groups of participants played the game and how successful their decisions were in terms of the game’s scoring system (by calculating the number of potential attacks defended out of a possible 33; and, by scoring how early in the game an investment was made). In our work we look not at the outcomes of decisions, but instead at the *process* by which the participants arrived at them, and the forms of risk thinking involved in the related discussion.

2.3 Data pipeline

A multi-phase, mixed method approach was used to analyse the data for each team:

Stage 1: The transcripts were manually coded to highlight where teams discussed risk. We used the following apriori codes, which were developed from the advice given on risk analysis in ISO270001 standard and NIST-800-53 [21, 22], and discussed and agreed upon by the authors: **Assets** Used wherever the teams have talked about a specific asset (e.g., the SCADA controller, the PCs or the databases) that is at risk.

Threats Used wherever the teams discuss potential actors who may be interested in attacking the organisation (e.g., criminals stealing data or saboteurs).

Vulnerabilities Used wherever teams have considered internal weaknesses to the current system (e.g., unsecured databases).

Impact Used wherever teams consider the potential consequences that could arise as a result of a particular action under consideration (e.g., fines from a data-breach).

The lead author did the coding on the dataset owing to their experience analysing decision-making data and familiarity running game sessions. The reliability of this initial coding was established by asking an independent reviewer (who is not one of the authors, but is an experienced security researcher) to code a random 20% subset of each transcript. In total, both reviewers coded 1293 paragraphs with a resultant Cohen's kappa of 0.71, demonstrating a good level of agreement on coding [1]. We are, therefore, confident in the quality of this initial coding.

Stage 2: Each round was then divided into 4 equally sized quartiles and the proportion of risk related discussion in each quartile was calculated. We used quartiles to help provide a greater level of granularity when assessing where risk conversation occurred during each round. We felt this was necessary because of the round-based nature of the game which may have encouraged teams to bookend their risk thinking during rounds. That is, teams may consider risk more when reflecting on the feedback received at the end of the last round; and they may also reflect on risk more toward the end of a round when they are finally committing to their investment decisions. By dividing each round into 4 equally sized quartiles we could assess this and compare between teams regardless of variation in the length of games. The top 5% of all quartiles (as measured by the proportion of all discussion coded) in each risk category were then selected to explore the process whereby risk thinking takes place. 33 quartiles were identified as containing a high proportion of risk thinking. A further random 10% of the other quartiles were also assessed to ensure that no areas of interest were excluded. We selected the quartiles with the highest proportion of coded discussion to direct analysis towards areas of the transcript which focused on risk thinking (rather than, e.g., meta-gaming, incidental discussion). We assessed a further 10% random sample to check that we did not miss any sections where teams talked about risk.

Stage 3: Graphs were created for the quartiles containing high levels of risk-discussion, using the process described in Section 2.4. We drew up the decision-making flow for the entire round containing any quartiles of interest, resulting in 19 distinct decision-making charts. These charts are available online as an appendix to the paper¹. The point of theoretical saturation was reached following four passes of the transcripts.

Stage 4: The 19 charts were then analysed to identify the core decision-making structures employed by the teams, as expressed in underlying regularities within the graph structures. In our findings, we discuss why these regularities are interesting and identify where teams use certain patterns more than others.

¹<https://github.com/benshreeve/Decision-making-charts>

“So if Mr Blue Head here clicks the link...”

Fig. 2. C11 Round 1

2.4 Visualising decision-making

The way that teams structure their conversations provides a valuable insight into the way that decisions are made. We borrow from speech-act theory in order to explore this process. At its most basic, speech-act theory suggests that people use language to perform actions i.e., to get a response of some kind from a listener [40]. For example, when someone asks a question they rarely have to declare that they are asking a question, instead the listener identifies that a question has been asked and provides a response. However, this isn't always the case. Sometimes a question can be asked and the listener instead offers up an alternative response—perhaps a counter-question, or maybe they miss that a question has been asked altogether. This relationship between speaker and listener helps us to explore how conversations develop. We have used the broad concept of speech-act theory, that is, the relationship between *illocutionary actions* (what the speaker is attempting to do) and *perlocutionary effect* (how the listener responds) to frame our analysis. This framing enables us to explore how the conversation develops, and in particular where understanding is developed, or in some cases breaks down.

The decision-making graphs were derived by analysing the transcripts of the rounds in detail using a speech-act lens to identify common decision-making features used by the teams. Figure 2 provides an example of this method and is the decision graph drawn up for team CI1's discussion and reasoning during Round 1. The graph nodes are numbered to show the sequence in which they occurred, whilst arrows show how discussion around one aspect related to the discussion of another. Note that teams did not always proceed between related considerations sequentially, in fact they would often refer back to points made a long way back in the conversation (either explicitly or by insinuation). Key quotations are included in graphs to help illustrate where teams employed deeper reasoning. Finally, the point at which a decision has been reached and an investment made is marked as a termination of the discussion.

The team in our example starts the game by identifying the lack of firewall as an immediate vulnerability (#1) with the potential for the lack of firewalls to result in a remote attack on the SCADA controller running the turbines (#2):

“So what immediately comes to mind is if there's no firewall then someone on the internet can just remotely connect to the controller.”

The team notes that firewalls are necessary on both sites because a single defended site might be vulnerable via the second undefended site. Teams then introduce a declared aim (#4):

“It would be interesting to know what we're fighting against”

This leads the team to consider whether to invest in the threat assessment or asset audit. During their discussion of which source of information to consult first they continue to acknowledge the need for firewalls on both sites. After some discussion the team decides to invest in the threat assessment (#8) the information from which then prompts them to consider network monitoring as a means of defending the server which they assume also presents a remote access target. The team refers back to the threat assessment, highlighting the threat of remote attacks as the more immediate priority. At this point the game master asks the team if they are referring to remote or physical attacks. The team then move on to consider the risk of physical attacks via espionage. They also briefly query whether they should do an asset audit. They think that industrial espionage would potentially aim to extract valuable data. One player suggests that antivirus may be a good counter to espionage but the other players do not think that a virus would affect the turbines:

“I think it's unlikely it will...I mean infecting us with a virus will do anything to the turbines.”

Team	Round	Nodes	Investment Options	Risk Thinking	Declared Strategy	Unavailable Choices	Declared Aims	Decisions	Team	Round	Nodes	Investment Options	Risk Thinking	Declared Strategy	Unavailable Choices	Declared Aims	Decisions
CA1	1	25	8	12	1	2	0	2	SA1	1	21	10	6	0	1	0	4
CA2	4	13	3	7	0	1	1	1	SA1	2	29	15	9	1	2	0	2
CI1	1	20	7	9	2	0	0	2	SA1	4	9	4	2	0	1	0	2
CI1	3	14	6	4	0	0	1	3	SA2	2	13	5	4	0	1	0	3
CI1	4	16	6	5	0	0	0	3	SA2	4	12	4	4	0	1	0	3
CI2	1	29	13	9	3	1	0	3	SI1	1	27	12	10	0	1	0	4
CI2	2	23	9	10	0	1	0	4	SI1	2	15	3	6	0	0	2	4
CI2	4	24	11	8	0	3	1	1	SI1	1	27	12	10	0	1	0	4
MI1	1	19	9	6	1	0	0	3	SI1	3	11	4	5	0	1	0	1
MI2	1	13	6	3	0	1	0	3									
MI2	4	2	1	0	0	0	0	1									

Fig. 3. Count of different graph nodes in each of the analysed rounds.

The notion of antivirus as a means of mitigating espionage are then considered (#15). Following this a player puts forward another strategy to have a firewall on one site plus antivirus (#17). The team then critiques this strategy, noting a further vulnerability of a malware attack via USB Thumb Drives before once again returning to the need for two firewalls or nothing:

“If we don’t get firewalls on both sites then they can always attack the other site and get access, but if we don’t get antivirus then the USB key scenario remains possible.”

The team debates this for some time before ultimately accepting that the implementation of firewalls on both sites is the biggest priority (#20).

Figure 3 shows the count of each type of node in the graphs we created from the round transcripts for all the teams. Discussion is dominated by talk about *game investment options* or *risk thinking* relating to vulnerabilities, threats, assets or the potential impact of a choice. Some teams made references back to prior investments or to investments that they wished were available. Teams sometimes employed mechanisms to help structure their decision-making such as *declared decision making strategies* or *declared broad security aims*. These mechanisms were used by teams to focus their discussion around a particular objective.

3 ANALYSIS

Figure 4 provides an overview of our core findings, starting with the initial coding of the transcripts to explore how risk thinking is distributed throughout teams’ conversations. This analysis is then used to select portions of the conversation (quartiles of each round) for further analysis where teams demonstrate high levels of risk thinking. These quartiles are analysed and visual representations (graphs) of the teams reasoning produced and analysed to reveal the underlying structure of cyber risk thinking (revealing *isolated*, *sequential*, *radial* and *complex thinking*), and reasoning approaches

Fig. 4. Framework showing how teams talk about risk, and different approaches to discussion.

used by teams (including *risk* and *opportunity-first* reasoning). The overall consistency and plasticity of the application of these characteristics is also derived.

3.1 Distribution of risk discussion

Figure 5 illustrates the proportion of discussion that each team devoted to each of the risk thinking codes (*assets*, *threats*, *vulnerabilities*, *impact*) across the course of the game. Each round is divided into 4 quartiles. The proportions reflect how much of the entire team's output in each quartile was so coded, and do not capture potential explanatory variables at an individual level (e.g., only one person talked about a certain category of risk), but only the overall team attention as expressed in discussion proportion. Note, team CA1 ran out of time playing the game and only completed 3 rounds of the game during their session.

These heat maps reveal a high level of inter-group variation in risk discussion practices. Contrast, for example, the plot for team CI2 with that for team MA1—CI2 considered all four types of risk thinking in all four rounds, while MA1 did not discuss risks during rounds 2 and 3, and never discussed risk in terms of *assets* or *impact*. Team MA1 use very little discussion during this period. Their decision-making is purely perfunctory in nature, for example, their reasoning for a core investment in the second round looks like this:

P1: We should go from this one?

P3: Yes I would say this is a CCTV surveillance in the plant. Network monitoring in the office.

No further reasoning is used; in those simple exchanges the team has decided what to invest in. Such differences in the depth and extent of reasoning used by teams potentially stem from the makeup of the teams and the way that different personalities evaluate risk.

In general, teams most often considered potential *vulnerabilities* during their risk discussion, this being more than twice as common as the discussion of *impact* (the next most common risk discussion category) or *threats*. Discussion of specific *assets* at risk was the least common, accounting for less than 1% of all risk discussion.

On average, the highest proportion of risk discussion takes place during the first round, followed by the second, and then the fourth—this pattern holds across all categories of risk discussion. One explanation for this pattern may be that risk discussion is front-loaded, as the teams come to

Fig. 5. Risk discussion heat maps of games, for all teams.

understand the game or game scenarios, producing a gradual decline which is then interrupted by the team’s awareness that they are entering the final round and have only one ‘play’ remaining. The emphasis on risk thinking during the first round in particular is indicative of teams making risk decisions under *unknowable uncertainty* where they have yet to encounter any attacks and as such have no experience of how the game penalises choices which could help to inform their decision making. This trend may be exacerbated by the round-based nature of the game and the way that unspent budgets roll-over between rounds—potentially affecting the way teams plan their responses, although we saw no evidence of this in the teams’ dialogue.

Considering all rounds at the quartile level, risk discussion is more commonly found in the first half (Quartiles 1 & 2) of any round. It may be that risk discussion is prompted by the feedback provided by the GM at the end of the preceding round, causing the teams to re-evaluate their position at the start of the following round. Teams are effectively moving from a position of *unknowable uncertainty* where they can only speculate about the type of attack they might encounter toward one of *known uncertainty*. That is, teams are aware of the attack(s) they have already suffered, meaning that they can work out which countermeasures to put in place to stop these attacks from occurring again. The described pattern holds for all subgroups except discussion of *impact*, where the fourth quartile appears just as important as the first two—possibly because decisions have to be made by the end of each round and so this is a natural location for the discussion of consequences. However, it is difficult to visualise this trend in the heatmaps because the impact code is relatively underused in comparison with the other codes.

Considering the games as a whole, the greatest overall concentration of risk discussion (in a single quartile) is during the first Quartile of the fourth round. This is particularly interesting given that the fourth round as a whole ranks only third for overall density of risk discussion. This suggests that this particular quartile is a critical time-point for the application of risk thinking, possibly because the first quartile follows the final instance of feedback from the GM. This quartile is also the most prominent overall for the risk discussion sub-types of *vulnerability*, *threats*, and *assets*. For discussion of risk related to *impact*, Round 1 Quartile 4 was the point with the highest concentration across the games. This may be due to apprehension becoming heightened as the teams approach the final decision point of the game and are unaware of what consequences they may encounter (again, teams are making decisions under increasingly *unknowable uncertainty*). They are, therefore, having to choose between a few items which they do not entirely understand. This results in greater use of risk thinking to evaluate the remaining options. For example, team CI2 spent round 4 attempting to negotiate between antivirus, CCTV and network monitoring:

P4: So in a way CCTV for the offices and possibly either network monitoring or the... sorry. CCTV for the plant and then either network monitoring or the other CCTV.

P2: So what don't you get? If you don't do antivirus you don't get protection, but we've been saying all along...

P4: If they've got good training they should have no reason to be installing it. So...

P2: And fully patched PCs now.

P3: We've got fully patched PCs. We've got employees that have been brilliantly trained.

P2: Yes. I think...

P1: Yes. I don't think we need an antivirus.

P2: ...encryption of the PCs. If you put CCTV in it, kind of, negates the need for encryption. Like X was saying about our philosophy, if you have secure premises then do you need encryption?

As a whole, these patterns suggest that most discussion of risk in game is primarily *reactive*—teams discuss risks immediately after the GM provides them with feedback about changes to the game state, and it would seem that they do this more extensively when they are aware this is the last piece of feedback they will receive. In these situations the nature of the game temporarily enables teams to make choices under *known uncertainty* rather than *unknowable uncertainty*. When playing the game, most of the discussion of risks came after the prompts providing new security information—implying that risk decision making, when playing the game, is primarily reactive. This parallels the way security threats are mitigated in the real world, such as the *penetrate-and-patch* approach [28]: until a risk is a clear and present threat, risk decision makers do not decide how to mitigate them.

In contrast to most risk dialogue, discussion of *impact*, in the game, tends to be *anticipatory*—teams discuss potential impact most of all when they are about to commit to a decision, and particularly when they are more uncertain about the potential consequences of a decision (at the start of the game). As such, it appears that *impact*, in the game, is a largely reflective form of risk thinking, and methods and approaches for understanding the impact of events, in the game, are most likely to be deployed later in decision-making conversations, *after* the range of vulnerabilities and threats have been considered and defined.

3.2 The structure of cyber-risk thinking

The 19 decision-making charts discussed in Section 2.3 were qualitatively analysed by two researchers using a content analysis-style approach to identify common mechanisms that teams

used to help structure their decision making (see Figure 6). These were then quantified to facilitate further comparison (see Table 3).

We identified four main mechanisms that teams use to structure their thinking: *isolated* (Figure 6a), *sequential* (Figure 6b), *radial* (Figure 6e) and *complex thinking* (Figure 6f)—these describe how a team’s discussion (and understanding) develops during the game.

These four mechanisms are informed by linguistic [40], management [24, 26, 49] and sociology [47] research.

We use the Speech-act theory (as discussed in Section 2.4) developed by Searle [40] to help us see where this sense-making occurs in the interactions of participants. We identify *Isolated* mechanisms by analysing the conversation to find where a participant performs an *illocutionary action* (e.g., asks a question) but does not received the anticipated response. That is, the anticipated *perlocutionary effect* doesn’t occur—for example, with an answer to the question, or perhaps a further question seeking clarification. Likewise, the *sequential* mechanism helps describe where a continuous string of these *illocutionary actions* and *perlocutionary effects* flow from one another in a continuous developing conversation.

The *radial* and *complex* mechanisms are derived from aspects of sense-making theory. Sense-making represents a key aspect of decision-making—it describes the process by which we “structure the unknown” by synthesising past experiences, known information and hypothesis to help the sense-maker “comprehend, understand, explain, attribute, extrapolate and predict” the phenomenon in question [44]. In our data set, teams use sense-making as they begin to develop their understanding of the game and the choices available to them. Management research often focuses on *sense-making* to help analyse the decisions made by individuals, teams and organisations during major disasters [24, 26, 49]. Suchman [47] describes a similar phenomenon when she introduces *situated actions*, using the analogy of Kayaking to help explain how an individual may stand on the side of a set of rapids and plan a route through them, but that once in the Kayak and tackling the rapids the actual route taken may vary wildly.

By combining these varied yet overlapping disciplines we are able to explore the decision-making used by these teams as they play the game and spot examples of these four mechanisms.

3.2.1 Isolated thinking. Isolated thinking (Figure 6a) describes situations where teams discuss one topic and then switch to talk about another without there being any explicit or implicit link. For example, in Figure 6c, team SA1 talk about the need for software patching and then move on to talk about the need for antivirus. We might have expected to see some form of link—perhaps one of the participants would say “*I think that software patches are important, but I feel that antivirus is more important at this point in the game*”—however, the participants do not explicitly make these links. Whilst they may have considered the link, they do not explain their conclusions to their team mates. Perhaps they considered it obvious, or unnecessary to explicitly discuss, but alternatively they did not consider it at all.

Nearly all of the teams from which these 19 graphs have been drawn demonstrate isolated thinking at some point (see Table 3). However, it appears that most teams do not use isolated thinking during Round 4. Since isolated thinking represents a disconnect (of some kind) in the conversation, it seems likely that by the time teams have got to the fourth round they have developed both a good understanding of the game itself and of collaborating with each other. Both of these in tandem could explain the lower prevalence of isolated thinking during the fourth round of the game.

3.2.2 Sequential thinking. Sequential thinking describes the process by which discussion of one idea leads logically into the discussion of another. For example, in Figure 6d, team SA2 talk about a prior attack they have suffered (#3) which leads them to talk about the fact that the attack

(a) Isolated thinking—That is, consideration of stimulus in isolation. Consideration of one does not prompt/is not related to the consideration of the next. Instead, stimulus 1 is considered and then stimulus 2 with no explicit or implicit overlap.

(c) Isolated thinking example from SA1, Round 1

(e) Radial thinking—These sequences involve using one core question or stimulus and then generating multiple related ideas from it.

(g) Radial thinking example from CA1, Round 1

(b) Sequential thinking—In this instance there are explicit or implicit links between one stimulus and another. Discussion of one prompts discussion of the next. This appears as a form of joined-up thinking.

(d) Sequential thinking example from SA2, Round 4

(f) Complex thinking—This is a more complex form of sequential think where dialogue cross-references and returns to other prior aspects of the conversation.

(h) Complex thinking example from SI1, Round 2

Fig. 6. Conversation development

Table 3. An indication of the level to which teams use different risk thinking approaches during 19 rounds selected for demonstrating high-levels of risk-related dialogue. ○ indicates a low-level of thinking about this topic in this round, ◐ indicates a medium-level, ● a high-level and a blank space indicates that the team did not demonstrate this kind of thinking in this round. We contrast this with the team’s overall score playing the game (taken from Frey et al. [16]).

Team	Round	Isolated Thinking	Sequential Thinking	Radial Thinking	Complex Thinking	Risk-based Thinking	Opportunity-based Thinking	Score (Reported in [16])
CA1	1	○	◐	●		◐	◐	27
CA2	4		◐		●	◐	○	27
CI1	1		●	◐		●	○	28
CI1	3	◐	◐	○	○			"
CI1	4	○	○	○		○	○	"
CI2	1	◐	◐	◐	○		●	31
CI2	2	●	○	○		○	○	"
CI2	4		●	○		○	◐	"
MI1	1	●	○			○		29
MI2	1	○	◐			○	○	28
MI2	4							"
SA1	1	●	○			○	○	22
SA1	2	●	◐	○		◐	◐	"
SA1	4	◐	○	○		○		"
SA2	2	○	●			◐	○	27
SA2	4		●				◐	"
SI1	1	●	◐			○	○	26
SI1	2	◐	●		◐	○	○	"
SI1	3	○	○			○	○	"

should have been neutralised by the anti-virus purchased in an earlier round (#4). Later on in the conversation the team talk about the need for PC encryption (#7) and refer back to the data exfiltration attack they had suffered (#3). Discussion of one item leads to the discussion of another, typically resulting in the exploration of related factors and as a result a deeper understanding is developed.

Sequential thinking is the most common form of thinking displayed by the teams (see Table 3) and is used by all teams during all rounds in the game. This is what we would expect to see given that teams are addressing a collaborative task and therefore are sharing information and ideas in order to identify their choices. However, there is a risk that teams may over-rely on sequential thinking

and end up following this logical train of thought to the exclusion of equally valid non-sequential options. For example, a team may spend all their time considering the investment in a firewall without even considering investing in antivirus. Frey et al [16] noted that one team in particular from this data set was prone to tunnel vision—we explore this in more detail in Section 5.

3.2.3 Radial thinking. Radial thinking represents a variant of sequential thinking whereby a number of ideas are generated in response to a single stimuli. This mode of thinking represents the commonly used *brainstorming* model of group idea generation, where teams take stock of where they are, and consider a range of options before deciding how to proceed. For example, team CA1 in Figure 6g stop in the third quartile and take time to consider their options:

“The next thing we’d want to do given that list of attacks would be to say *well, what are the consequences of each one?* And, *what’s the cost to us of each one?* And then and what the probability is for each one and make our decision based on that.”

They do just that, considering their various options before electing to install a firewall, having considered each threat in turn.

Few teams demonstrated radial thinking, with only CA1 making an express point of brainstorming risks. The remaining teams demonstrated radial thinking in a limited form, only ever exploring 2–3 options at a time. Given how widely brainstorming is used, we might have expected more teams to demonstrate more radial thinking.

3.2.4 Complex thinking. This is the most advanced form of discussion, where teams not only develop thoughts sequentially but cross reference prior ideas in order to explore them in greater depth. Complex thinking differs from radial thinking by referencing past thoughts or discussion points and considering how these would be impacted or affect the current trail of thought. This act of reflection is important as it lets the teams return to old ideas and consider whether they are still valid or have gained new significance. An example of this is in SI1 Round 2 (Figure 6h): the team consider the possibility that there is a threat from unpatched controller firmware in the SCADA system. Later on they think about the difficulty involved in exploiting the controller’s firmware and developing an exploit. They reason that it requires experience, so the attack is less likely (than some of the other potential attacks), but also then note the possibility that a compromised controller could be used to attack other devices. They reflect that, in their experience, they haven’t seen this often and that typically attacks (of this nature) often originate from PCs rather than unpatched controllers, and so they do not opt to upgrade the controllers.

This sort of complex thinking, including returning to past suggestions with more considered analysis, is often taught as a key component of good decision making and yet we very few teams demonstrated it. Only two of the 19 rounds analysed demonstrated this form of thinking—and minimally so at that.

In summary, our analysis finds that, in the teams we studied, teams were most likely to use *isolated* and *sequential* mechanisms to help structure their cyber-decision making. Teams were more than twice as likely to use these mechanisms than *radial* thinking approaches and more than three times as likely to use them as *complex* forms of thinking. Such an extreme difference in usage is particularly noteworthy given the emphasis that is placed on the use of structured risk assessment techniques within Cyber security [21, 22]—the teams we studied appeared not to use these techniques. Teams were slightly more likely to use *sequential* thinking than *isolated*, this is hardly unexpected given that teams are addressing a collaborative problem-solving exercise.

3.3 Reasoning Approaches

Alongside the mechanisms that teams use to structure their conversation we have identified two key forms of reasoning utilised by teams: *risk-first* and *opportunity-first*. These represent two different styles of reasoning observed in the game. We make no claim as to which approach is better, but the fact that these two approaches are seen is notable. We note (in Figure 8) that teams in the game use both forms of reasoning. This suggests that these two forms of reasoning are not mutually exclusive, but rather that they represent different approaches participants take whilst playing D-D—and that these two approaches to decision making might be seen in other decision making scenarios with limited choices and information.

3.3.1 Risk-first. Risk-first reasoning refers to the case where discussion is initiated by consideration of a risk. Teams start by identifying a risk which then prompts them to explore the means of negating the risk. In Figure 7c the participants begin by thinking about the possibility of a staff member making an error and how it could lead to the threat of a malware attack. They say:

“So if Mr Blue Head here clicks the link and gets some malware, does the antivirus detect that?”

They have identified a threat, and a possible attack vector, and their risk-first reasoning has led them to consider a possible mitigation strategy: the antivirus investment option.

This risk-first approach enables teams to identify vulnerabilities in their current systems and then consider which threat actors are most likely to exploit these—and then identify the optimum defence. Almost all of the rounds analysed demonstrated some risk-first reasoning (see Table 3). However, most teams only use a small amount of risk-first reasoning when playing. This would appear to contradict the majority of the existing literature which suggests that a defensive stance should be built on a risk-centric analysis ([21, 22]). In practice, the teams we studied tended to demonstrate only a passing focus on traditional risk analysis.

This is potentially explained by the fact that a risk-first approach is entirely reliant on the ability of the decision-maker to identify risks. Whilst it may be possible to gauge the extent to which vulnerabilities exist within a system—such as by measuring the rate at which patches are applied—it is much more difficult to ascertain who the main threat actors are and their potential attack vectors. This risk-first approach represents the tradition within management, but is not necessarily the best approach for all organisations in terms of cyber security. It is interesting to note that, for the majority of teams, initial game choices are consistent with this—choosing to play the *threat assessment* card first—seeking to understand the threat landscape ahead of anything else. This is despite teams recognising post-game that ultimately the external risk landscape is (by comparison with the internal vulnerability landscape) relatively unknowable. And yet despite this teams continue to place a greater emphasis on understanding the threat actor and vectors which for the most part they can only speculate about ahead of reviewing the actual tangible assets which they have to identify known vulnerabilities.

3.3.2 Opportunity-first. Opportunity-first reasoning, by contrast, begins with the identification of investments or opportunities before then considering what risks are associated with these possible choices. In Figure 7d for example the team start by exploring whether they should use their funding to invest in one of the encryption investments. They start with the investment and then note that actually there is a related risk (out of date operating system) which would negate the investment:

“There’s no point encrypting it if they can pull it out anyway. It’s running Windows XP.”

(a) Risk-first—Teams start off by identifying a risk/risks and then seek to identify the optimum mitigation.

(b) Opportunity-first—Teams begin by exploring the investment or opportunity available and then seek to evaluate its effectiveness by considering the potential risks that could be mitigated by it/raised as a consequence of investing in it.

(c) Risk-first example from CA1, Round 1

(d) Opportunity-first example from CA1, Round 1

Fig. 7. Forms of reasoning

The nature of D-D means that teams are always making decisions about cards (investment options) that are placed in front of them. The design of the game, therefore, encourages opportunity-first reasoning as teams collaborate to identify which investments to make from a finite selection of known options. Despite this, teams playing the game used both opportunity-first and risk-first reasoning equally. One possible explanation for this is that, within the context of the game, both approaches tend to result in the same final decision. For example, a team which utilises a *risk-first* approach identifies which particular attack vector they think is most likely (perhaps they think a denial of service attack is a major threat). This team then has to go ahead and work out which investment will provide them with the best protection against this attack, ultimately deciding that on the firewall. By contrast a team which uses an *opportunity-first* approach may be considering the firewall alongside multiple other options and then decide that the threat of a denial of service attack against their organisation is the current biggest problem and so invest in the firewall. Both teams end up at the same decision but both have arrived via different routes.

The key difference between opportunity-first and the risk-first reasoning is that the list of investment opportunities are inherently knowable (in that they exist as tangible cards). This extends to the real world where it is relatively straightforward to perform a standard accredited audit of a company's infrastructure and identify opportunities for improvement. By contrast, it is much harder to identify who is likely to be interested in attacking an organisation, why, and by what means. The effort and effectiveness of these two approaches is also likely to be directly linked to the scale of the organisation involved. A small organisation is going to find it much easier to identify their vulnerabilities than a large one, while a large organisation may well have a history of attack information to fall back upon, enabling them to profile their attackers using resources that a small organisation would struggle to match.

Figure 8 provides a comparison of how the use of these critical thinking approaches changes over time. It is important to note that most teams vary in the reasoning that they use as the game

Fig. 8. Comparison of critical thinking across rounds

develops, and that across all 19 portions, teams changed their strategy as the game progressed. Both management teams (the *M* teams) started the early rounds using a *risk motivated* approach—similar to threat modelling, whereas the security practitioners (the *S* teams) started without any discernible approach but over time moved to using a threat and opportunity based approach. This may be a sign that managers are more familiar with the concept of threat modelling—how we would hope they approach the task—and they therefore apply this approach from the start, whereas the practitioners wait to see how the game unfolds and then chose their approach accordingly. The computer scientist teams we examined (the *C* teams) all started their games incorporating some level of opportunity motivated thinking. This is notable as these teams typically performed better than other teams—with Team CI2 doing the best of all (Table 3), perhaps suggesting that an initial careful consideration of the available mitigations in the game may lead to better decision making overall.

3.4 Application of risk thinking

Our analysis revealed further differences in the way that teams applied structuring and reasoning approaches: *consistency* and *plasticity*. *Consistency* helps to describe the way that teams decision-making changed between rounds. For example, team SI1 (see Table 3) were one of the most consistent teams—in all three rounds analysed they demonstrated similar levels of isolated and sequential thinking and risk-first and opportunity-first reasoning. *Plasticity* meanwhile helps to describe how readily teams are able to switch between the different structures of risk thinking during discussions. Whilst plasticity and consistency can help to describe how the decision-making of teams changes during the game we are not attempting to gauge the quality of decision-making. Any attempt to do so is problematic [11, 27]. Further research will be necessary to show if consistency and plasticity of decision-making exist outside of the D-D context.

3.4.1 Risk thinking plasticity. One of the key findings that emerges is the relationship between the quantity and depth of dialogue, and with overall quality of decision making. Team CI2 (who

Frey et al. [16] highlight as the highest scoring team) utilise the most overall decision nodes (see Figure 3; 76 nodes in total) and exhibit the highest number of characteristics (14 in total across three rounds, see Table 3). They also demonstrate a much greater use of these characteristics than the other teams reported herein with an average depth of 2.67 (where the levels on Table 3 are given the following values: no thinking shown=0, low-level=1, medium-level=2 and high-level=3).

Plasticity describes the ability of teams to switch between decision making approaches. It helps to explain why team SA1 utilise a similar range of characteristics (13 in total across three rounds, see Table 3) and to a depth which is only slightly less than that of team CI2 (averaging a depth of 2.34) and yet, Frey et al. [16] consider them to have played the worst game of all 12 teams. Teams SA1 and CI2 demonstrate different levels of plasticity which explains the difference in in-game performance. For example, team SA1 spend a lot of time in each round performing isolated thinking, team CI2 by comparison utilise a wider range of characteristics to a medium-high level. Future work should examine whether plasticity affects real-world decision making similarly.

3.4.2 Consistency of risk thinking communication. Closely related to the plasticity of teams is the consistency with which teams apply risk thinking. Teams CA1, CA2 and MI1 for example, were only seen communicating risk thinking in a single round, whereas other teams in Table 3 demonstrated risk thinking through discussion in multiple rounds. Referring back to Frey et al. [16], teams CA1 and CA2 represent the lowest scoring teams comprised of Computer Scientists—this may give a potential sign that consistently communicating risk between team members is a reason for scoring well in the game. Future work should check whether teams of cyber decision-makers make more effective decisions when reacting to threats in the *real world* when proactively communicating risk thinking between team members.

4 RELATED WORK

Existing research on cyber security risk decision-making largely falls into five categories: (i) methods and tools for supporting risk decision-making; (ii) risk communication across an organisation; (iii) how individuals perceive and evaluate cyber security risks; (iv) serious gaming, and (v) decision-making during serious gaming. We discuss each of these next and how our work complements this body of research.

4.1 Methods and tools for supporting risk decision-making

One of the dominant areas of research is around methods for assisting cyber security decision-making. This research usually relates to quantifying cyber-risk to help inform decision-making. Some research also takes financial factors into account to try to give an indication of the potential financial ramifications for decisions.

Early work on risk analysis by Rainer Jr et al. [37] suggests that managers responsible for IT security typically only utilise a single risk methodology, when instead they would gain better insight by incorporating a range of different risk assessment methodologies. They note in particular the need for some empirical basis for cyber risk.

The quantification of risk has become increasingly popular within organisations, with formulae like $\text{Risk} = \text{Threat} \times \text{Vulnerability} \times \text{Consequences}$ or metrics such as CVSS [51] becoming synonymous with cyber-risk evaluation. The issues associated with such approaches are well known [10]—namely that these formulas have a tendency to over-simplify the risk assessment process, and often fail to account for attackers changing their strategies.

Some 17 years after Rainer Jr et al. [37] and the search for effective empirical representations of risk continues. For example, Bodin et al. [5] propose a metric—the *perceived composite risk*—as a mechanism to allow risk decision makers to combine multiple risk-metrics into a single value.

They noted that in a cyber security context *risk* is poorly defined and that those responsible at board-level within companies needed to better define the term in order to manage it more effectively. Earlier work by Bodin et al. looked at creating a different model to evaluate different mitigation strategies for dealing with risks by placing a *dollar-value* on it, based on the danger to confidentiality, integrity and availability [4]. This work reinforces that of Rainer et al. [37], concluding that cyber risk decision makers should consider a greater range of threats and impacts.

Closely related to this discussion is the tendency for organisations to use financial evaluation as a core part of decision-making, and cyber security is no different. This results in work such as the Annual Loss Expectancy (ALE) model [33] which attempts to forecast the expected impact of a range of threats occurring within a 12 month period. ALE is used as the basis for a range of security investment decision-support approaches. It is closely related to the Return on Security Investment (ROSI) model (e.g. [7]) which attempts to provide an indication of the benefits gained from security investments. Such approaches are heavily reliant on the quality of the risk evaluation. Similarly, work by Gordon and Loeb [20] proposes an economic model to determine the optimal amount to invest in security to offset the risk against a particular data-set. Their approach assesses risk in terms of offsetting the vulnerabilities of a particular data-set with the potential impact if it were stolen.

Recent work by Moore et al. [30] meanwhile, report on a survey to explore which factors influence executives’ cyber security investment decisions. They suggest that executives often use frameworks (e.g., NIST or ISO) to help structure their decision making. They suggest that most executives are concerned with process measures—that is, how gaps in an organisations existing setup can be found and fixed—rather than on the potential consequences of choices. This finding is contrary to the aforementioned research which promotes investment based on potential impact.

Our work is not about suggesting a model for making decisions that can be implemented within an organisation. Instead we study the decision making processes that people use *in practice* to make risk decisions, rather than the frameworks and models they might be trained to use.

4.2 Communicating risks

Communication is an important aspect of risk decision-making and response and there is a range of work related to it, including much of the work mentioned on sense-making [24, 26, 49]. There is a growing field of work that explores risk communication within a cyber security specific context. For example, Feledi et al. [15] developed an information security ontology to help companies to share information and describe a shared vocabulary [15]. Their ontology described assets, threats and controls and was specifically designed to help risk-management and compliance tasks—it does not, however, mention risk as a metric for the likelihood of an attack, instead talking about vulnerabilities and their severity. This aligns closely with the apriori codes we have used for our initial coding (as described in Section 2.3) where we have coded to highlight where teams are talking about *assets*, *threats*, *vulnerabilities* and [potential] *impact* [of decisions].

Coles-Kemp and Overill noted that businesses often have poor mechanisms for communicating security risks, with assessors creating findings that the business cannot understand [9]. This of results in scenarios where security risk assessments become a box ticking exercise—done for the sake of doing a risk assessment rather than for the sake of securing the company’s assets. They argued for the need for a cyber risk facilitator to help explain the dangers from cyber risks, but noted that standards such as ISO 27001 did not mention the role.

In our work we find that some teams talk about risks and dangers while playing the game, but that those who do often do so in an unstructured manner—this suggests that risk thinking is not just a box ticking exercise but potentially something that people don’t think about at all without specific prompting arising from professional or organisational protocols.

Other work considers the impact that trust has on the communication of risk, with Nurse et al. [34] providing a literature review. They found that many different factors that impacted how trustworthy people found risk information [17, 25, 32, 36], but that there had been a lack of research in *usable* cyber risk communication, and that numerical approaches may not be a suitable approach for all cyber risk managers. As previously stated, we found in our study, that no participants used numerical metrics when discussing risks.

4.3 Cyber-risk perception

Risk perception is closely related to much of the work in Sections 4.1 and 4.2, it differs in that the emphasis is on understanding the socio-cognitive processes involved in identifying and evaluating risks. There are some obvious overlaps with the management and sociology research explored in Section 3.2. Much of the early work is not from a general risk background (e.g., [38, 39, 42] and is included here for context.

Renn looked at risk perceptions in organisations from a social science perspective [38], observing that risk decision makers are under considerable political pressure from the public as well as from experts. He noted that, when the general public were asked what they would like risk decision makers to do, the public don't only want the decision makers to take actions that reduce risks, but to also perform the actions that they felt the decision makers *ought* to do—even if they didn't actually reduce the risk. In our study we also saw these patterns when playing the game; with several teams opting to start the game playing the firewall as a “*no-brainer*” without actually considering what it was they were protecting.

In his later work, Renn [39] goes on to suggest that risk assessment involves three key aspects: (i) identification of, and where possible estimation of hazard; (ii) assessment of vulnerability and/or exposure; and (iii) an overall estimation of risk combining the likelihood and severity. In our analysis all three of these aspects form part of the *risk-first* and *opportunity-first* reasoning that the teams show.

There is a range of work that explores cyber-risk perception. Recent work by Stevens et al. [45] has explored whether training staff in the New York City Cyber Command to use Centre of Gravity threat modelling would improve their cyber security decision making. They found that staff who had completed the training were better able to spot new cyber security threats and address them. Their ability to perceive risks had been improved by training them to consider the wider implications of choices they were making and data they were evaluating.

M'manga et al. [29] meanwhile explore how security analysts improvise, combining aspects from a range of risk analysis methods in their work. They refer to this combination as *folk risk analysis*. They found four groups of factors that influence the analysts interpretation of risk: *awareness*, *communication*, *tool capabilities* and *individual capabilities*.

Jalali et al. [23] have explored whether experts or non experts are any better at handling uncertainty in predicting cyber incidents or understand delays when implementing cyber security capabilities. Their work suggests that both groups had issues understanding delays in the implementation of capability, and that both groups exhibited similar errors when dealing with uncertainty of cyber incidents. As part of their work they specifically call for training of decision makers and for further research into mental biases in cyber security. The work in our paper begins to address this, by providing some insight into the way that teams go about making decisions during the game. However, further research is needed to explore whether these decision-making mechanisms exist outside of the D-D context.

Work by Downs et al. [14] uses a mental models approach in order to explore how inexperienced users make decisions about phishing emails; namely if/how they identify such emails and how they respond to them. Their later work [13] then explores the underlying behaviours that lead people to

fall for phishing emails. In particular, the way that people perceive the risk of phishing emails and potential consequences. Their focus, however, is on individual decision making, in this paper we move beyond this to explore how collective decision making occurs.

Other research explores more holistic approaches to risk evaluation. The work by Frey et al. [16] highlights that, although inexperienced in cyber security, managers can still make sensible investment choices because of the way that they consider risk. Related to this is the work by Straub et al. [46], who propose a security risk planning model consisting of five stages: *Recognition of Security Problems, Risk Analysis, Alternatives Generation, Decisions and Implementation*. This suggests a chain of risk decisions at each step where participants brainstorm possible problems and alternatives.

Unlike Stevens et al. [45] we have not set out to explore if we can affect risk-perception and decision-making of teams. We are instead interested in exploring the way that teams perceive cyber-risk within the context of D-D. In this way our work is more closely aligned with that of M’anga et al. [29] and Downs et al. [14].

4.4 Cyber security games

Games have become a popular way to raise awareness of cyber security [2, 3, 12, 16, 19, 41]. Many of these games are designed for educating security students (e.g., [3, 19, 31]). D-D, by contrast, sets out to provide a simplified representation of real-world decisions within an industrial organisation. No prior knowledge is necessary, and the game is not designed with the express purpose of teaching, but rather to provide a conduit for people to demonstrate the decision-making they might have used if working in the real world.

Other games and related activities require participants to have experience or technical skills, for example, Bock et al. [3] who describe the development and use of a king-of-the-hill style competition to help encourage their students to get more involved in cyber security.

The work by Beckers and Pape [2] is perhaps the closest to D-D. Their game challenges participants to help extract security requirements to help an organisation defend against social engineering attacks by helping the participants better understand the attacks and methods used [2]. Their game encouraged the participants to consider the risks to a company, develop attacks and then rate them by their plausibility. They did not, however, look further into how the participants came up with and reasoned about risks associated with the attacks.

4.5 Decision-making in games

There is a further subset of research that explores decision making in games specifically (as opposed to the broader work on decision-making and sense-making discussed in Section 3.2).

Bornstein et al. [6] for example, use the Centipede game to explore whether individuals or groups were primarily concerned with winning the game—this means that participants playing against each other are only concerned with maximising their own payoff during the game. Their analysis suggests that neither individuals nor groups were fully motivated by the desire to maximise payout during the game. However, they did find that groups were slightly more inclined to pursue greatest return. Of the D-D teams studied in this paper, we have found no situations where teams exhibited similar goal-driven motivations.

Work by Xu et al. [52] suggests that there are five categories of social interaction that occur during board games: *Reflection on gameplay, Strategies, Out-of-game, Game itself* and *Chores*. *Reflection on gameplay* describes the way that teams react to and reflect on gameplay following a move. *Strategies* relate to the way that teams decide how to play before making a move. We have identified these in our analysis as *risk-first* and *opportunity-first* reasoning. *Out-of-game* interactions are those that occur but that not directly related to the game. They note that often the

game itself forms the centre of the interactions, that is, commenting on and reacting to the game itself, as an artefact in its own right. Finally, they have identified a subset of interactions that they term *chores* which form a key part of gameplay. These are interactions such as rule learning, waiting, moving physical parts of the game. This process is readily apparent in our analysis, with teams dedicating a lot of time during the first round toward understanding how the game works.

Work by Gladstein and Reilly [18] explores how group decision-making changes when a group is under threat from external factors. Their work suggests that when a group faces a change in circumstance or external threat then their ability to process information is restricted and teams tend to constrict control. This results in rigidity in response. This break down of reasoning is evident in our analysis of teams playing D-D where an illocutionary action fails to be followed by the expected perlocutionary effect—in our visualisations of decision-making this occurs where a sequence of reasoning is disrupted.

St. Germain and Tenenbaum [43] have explored the decision-making processes of expert and non-expert poker players by asking players to think aloud whilst playing through hands. They found that expert and intermediate players outperformed novice players in terms of decision-making performance. They also found that expert and intermediate players were capable of processing far more cues as part of their decision-making. Their experience meant that they could consider a wider-range of possible solutions. This finding in particular relates closely with Weick's [50] work where he noted that people extract cues from the context of a situation to better understand what is going on.

We are not therefore the first to consider exploring the decision-making of people playing games. This work is one of the first to explore decision-making in the context of a cyber security decision-making game. Our work differs from the works described here in a range of ways, unlike Bornstein et al. [6] we have not set out to explore a specific hypothesis. We focus on understanding how the decision-making naturally develops during D-D. The broad categories of interactions identified by Xu et al. [52] are insightful, but remain very high level. Our findings explore the decision-making trends in terms of linguistic interactions which enables us to identify more specific decision-making characteristics. Our work has some overlap with that of Gladstein and Reilly [18], in that we have identified mechanisms that are perhaps associated with the breakdown of information processing.

5 DISCUSSION

5.1 Risk thinking and existing decision-making frameworks

Existing frameworks and advice for risk thinking tend to advocate a risk-first approach to problem solving. For example, the UK Cabinet Office Risk Thinking Model Office [35] proposes a simple iterative process starting with the identification of risk followed by mitigation:

Identify risks → Assess Risks → Build resilience → Evaluate resilience

However, our analysis suggests that, at least in the teams we studied, teams actually use other approaches for evaluating risk, e.g., opportunity-first reasoning. It seems likely, therefore, that there are other viable approaches. Indeed within our study team CI2 which performed the best overall (see Frey et al. [16]) made much more use of opportunity-first reasoning than risk-first.

We would posit that there are a number of assumptions that have to be made in order to use a risk-first approach which aren't necessary for opportunity-first reasoning. Foremost of these is the assumption that people can actually identify cyber risks. This is especially problematic for non-technical individuals where it is difficult to consider knock-on consequences. These risks are therefore unknown and even unknowable in some cases. By comparison teams are much more likely to be able to identify more tangible opportunities such as investment in defences and training.

We, therefore, suggest that future risk thinking methodologies take this into account—there is at least one more valid approach to critical thinking available.

Related to this is a tendency to frame cyber risk analysis in the same way that risk is calculated in engineering and finance fields. Such approaches fail to take into account the inherent complexity when considering cyber-risk. In traditional engineering, risk is often binary in nature. One is able to evaluate the likelihood of a material failure through testing and then go on to estimate the likely knock-on impact of a failure. However, the inter-connected nature of computing and communications systems makes this impossible to resolve for cyber security. There are multiple points of failure which are ultimately unknowable. Yes, one can make some good guesses but there will always be new compromises to be found. The knock-on effect of any of these points of failure being exploited represent a further issue, a major failure may occur and be entirely recoverable or it may destroy an organisation. Such complexity means that decision-makers are forever *dancing in the dark*, treading a line between the unknown and the unknowable and attempting to cover for both.

5.2 Risk thinking and decision quality

We refer back to the original paper by Frey et al. [16] in order to explore the relationship between the risk thinking characteristics reported in this paper and quality of decision-making. In terms of overall performance, there appears to be no clear link between the scores reported in Frey et al.’s. [16] original paper and the 19 rounds (8 distinct teams) selected as demonstrating high levels of risk thinking dialogue. Both the highest scoring team (CI2) and the lowest scoring team (SA1) demonstrate risk thinking and yet the outcome of their dialogue is very different. Nevertheless, there are still some interesting findings. For example, team CI2, which has the most balanced approach (demonstrating risk thinking in 14 of a possible 18 instances) in terms of structure and reasoning, were the best performing team in the original performance analysis [16]. Team CI2 were “surprised by their excellent result, as they were constantly expecting a disaster to happen until the very end” [16] which perhaps explains why they employed such a mix of approaches, demonstrating more rich thinking in round 4 than most teams. By contrast, Team SA1 consistently demonstrated isolated thinking and used relatively little sequential thinking and were the worst performing team according to Frey et al. [16] and suffered from “tunnel vision”. Team SI1 were the second-worst scoring team in the original analysis, also demonstrating tunnel vision by neglecting to consider the likelihood of data exfiltration attacks. In our analysis the team used both isolated and sequential thinking (and even some complex thinking) but used relatively little reasoning.

5.3 Risk thinking and group composition

The original paper [16] reports the experience of participants as an important factor. Teams CA2, MA1, MA2, MI1 and MI2 in particular are highlighted as lacking either (general) experience or technical background (sometimes both). Frey et al. [16] note that this lack of experience often manifests itself in very limited reasoning. In our analysis teams CA2, MI1 and MI2 demonstrated risk thinking dialogue in only a single round. MA1 and MA2 failed to demonstrate a significant level of risk thinking and so weren’t even included in the 19 quartiles of interest explored in more depth. MI2 round 4 is an especially interest anomaly, the team apparently demonstrated risk thinking to the extent that they were one of the 19 quartiles selected for further analysis. However, the resulting graph of their reasoning is so basic that they actually fail to demonstrate any of the six highlighted risk-thinking characteristics.

In comparison teams SI1, SI2, CA1, CI1 and CI2 were all reported as experienced in the original paper [16] using richer anecdotes to explain their reasoning. Of these, four teams demonstrated risk thinking according to our analysis in order to be selected for further analysis (CA1, CI1, CI2 and SI1)

with the latter three teams exhibiting risk thinking across multiple rounds. It therefore appears that there is a relationship between experience and consistent application of risk thinking. CA2 round 4, represents another interesting variation, providing the closest example of fully-featured complex thinking that we identified in the transcripts. This seems to suggest that perhaps these teams are capable of applying risk thinking, but lack the confidence to do so. Future research should look to explore this further working with a larger data consisting of both established and newly-formed teams.

5.4 Visualising risk thinking

By visualising the collective risk thinking process we have been able to explore the underlying socio-cognitive processes used by teams when making cyber security decisions. The graphs created have enabled us to identify two constituent parts of cyber security risk thinking: firstly, the structure of the conversation and second, the application of risk-first or opportunity-first reasoning. These two parts often occur concurrently but perform two distinct functions, the structure of the conversation describes how teams develop their reasoning through their dialogue. The reasoning aspect then describes the two main approaches that teams demonstrate for evaluating decisions.

Both the method of visualising and identification of these risk thinking characteristics represent important findings. The visualisation enables us to analyse patterns that we known exist in conversations in a new way. For example, it is interesting to note how few teams demonstrate radial thinking, given that brainstorming is arguably the most commonly used problem solving method used extensively in teaching from a young age and within organisations. Using this approach of mapping socio-cognitive interactions we have been able to identify the distinction between conversation structure and reasoning.

5.5 Threats to validity

There are a number of threats to validity relating to our method that we must acknowledge. Internal threats to validity include the coding of the Frey et al. [16] transcripts; this was done by the lead author and then a second independent coder who reviewed a random 20% of the transcripts with a Cohen's kappa of 0.71, suggesting a good level of coding agreement [1]. There are further threats to validity that relate to the gathering of the original data-set, these are covered in more detail by Frey et al. [16].

The decision-making of the teams is also likely to be influenced by the composition of the teams and the depth to which they engage with the exercise. Future work should explore the relationship between team composition, quantity of dialogue and relative quality of decision-making.

In terms of external threats to validity, it is important to reiterate that the characteristics reported herein are derived from studying a limited sample of teams playing a game. The generalisability of these characteristics therefore needs to be established through careful ethnographic observation of teams making decisions in real-world situations. However, access to real world teams which will allow observation (and potential critique) of their decision-making is problematic. D-D therefore provides as effective a decision making scenario for us to explore cyber-risk decision making in teams as is currently possible. Furthermore, our findings are artefacts of the game itself. We do not claim that a person who plays a high scoring game will necessarily make excellent cyber security decisions. Our findings are about how people make decisions when playing a game—the game is an effective simulation for modelling how cyber-risk decisions are made when teams from different backgrounds make decisions [16]. This forms a basis to encourage teams to avoid *isolated* and *sequential* thinking when making cyber-risk decisions and consider a combination of reasoning approaches as reflected by *opportunity-first* and *risk-first* reasoning in our analysis. This can help

enable a fuller exploration of the risk through decision-making grounded in deeper risk-thinking and reasoning.

6 CONCLUSION

Our work has focused not on *what* decisions security decision makers make, but rather on *how* they make them. We identify several different patterns of risk thinking and gain insight into how their decision-making process changes as attacks progress and time passes.

There was little evidence of teams reflecting on past information, or *brainstorming* different possible solutions (as evidenced by the lack of *radial* and *complex* thinking patterns). Instead, the bulk of the observed decision-making process was either *isolated* or *sequential* thinking—with participants starting from either a threat or a possible mitigation investment, and then either discarding it or moving on to their next thought as if dancing in a conga-line. This may suggest that teams did not, in general, reflect or consider different solutions; but it may be the case that individuals in teams did not feel a need to communicate their thought processes to other team members. Future work should investigate this further.

Further work should look to explore why cyber risk practitioners are making the decisions as they currently are—are they trying to follow a methodology, or are they just fighting fires as the security events progress? How do these decision process take place under different scenarios—in our study the participants were making security risk decisions whilst playing the D-D game, which introduces new threats over 4 rounds: how would cyber risk decision makers make decisions if this occurred over a shorter more pressured timescale (say an attack to their databases currently occurring) compared to a longer period where they may not be aware of any breaches occurring? Our work gives us a framework for mapping how these decision makers are making their choices.

7 ACKNOWLEDGEMENTS

This work is supported in part by EPSRC Grants MUMBA: Multi-faceted Metrics for ICS Business Risk Analysis (EP/M002780/1) and DYPOSIT: Dynamic Policies for Shared Cyber-Physical Infrastructures under Attack (EP/N021657/2)

REFERENCES

- [1] Douglas G Altman. 1990. *Practical statistics for medical research*. CRC press.
- [2] Kristian Beckers and Sebastian Pape. 2016. A serious game for eliciting social engineering security requirements. In *2016 IEEE 24th International Requirements Engineering Conference (RE)*. IEEE, 16–25.
- [3] Kevin Bock, George Hughey, and Dave Levin. 2018. King of the Hill: A Novel Cybersecurity Competition for Teaching Penetration Testing. In *2018 USENIX Workshop on Advances in Security Education (ASE 18)*. USENIX Association, 9.
- [4] Lawrence D Bodin, Lawrence A Gordon, and Martin P Loeb. 2005. Evaluating information security investments using the analytic hierarchy process. *Commun. ACM* 48, 2 (2005), 78–83.
- [5] Lawrence D Bodin, Lawrence A Gordon, and Martin P Loeb. 2008. Information security and risk management. *Commun. ACM* 51, 4 (2008), 64.
- [6] Gary Bornstein, Tamar Kugler, and Anthony Ziegelmeyer. 2004. Individual and group decisions in the centipede game: Are groups more “rational” players? *Journal of Experimental Social Psychology* 40, 5 (2004), 599–605.
- [7] Huseyin Cavusoglu, Birendra Mishra, and Srinivasan Raghunathan. 2004. A model for evaluating IT security investments. *Commun. ACM* 47, 7 (2004), 87–92.
- [8] Clare Chua Chow and Rakesh K Sarin. 2002. Known, unknown, and unknowable uncertainties. *Theory and Decision* 52, 2 (2002), 127–138.
- [9] Lizzie Coles-Kemp and Richard E Overill. 2007. On the role of the facilitator in information security risk assessment. *Journal in Computer Virology* 3, 2 (2007), 143–148.
- [10] Louis Anthony Cox, Jr. 2008. Some limitations of “Risk = Threat × Vulnerability × Consequence” for risk analysis of terrorist attacks. *Risk Analysis: An International Journal* 28, 6 (2008), 1749–1761.
- [11] James W Dean Jr and Mark P Sharfman. 1996. Does decision process matter? A study of strategic decision-making effectiveness. *Academy of management journal* 39, 2 (1996), 368–392.

- [12] Tamara Denning, Adam Lerner, Adam Shostack, and Tadayoshi Kohno. 2013. Control-Alt-Hack: The Design and Evaluation of a Card Game for Computer Security Awareness and Education. In *CCS*. ACM, 915–928. <https://doi.org/10.1145/2508859.2516753>
- [13] Julie S Downs, Mandy Holbrook, and Lorrie Faith Cranor. 2007. Behavioral response to phishing risk. In *Proceedings of the anti-phishing working groups 2nd annual eCrime researchers summit*. ACM, 37–44.
- [14] Julie S Downs, Mandy B Holbrook, and Lorrie Faith Cranor. 2006. Decision strategies and susceptibility to phishing. In *Proceedings of the second symposium on Usable privacy and security*. ACM, 79–90.
- [15] Daniel Feledi, Stefan Fenz, and Lukas Lechner. 2013. Toward web-based information security knowledge sharing. *Information security technical report* 17, 4 (2013), 199–209.
- [16] Sylvain Frey, Awais Rashid, Pauline Anthonysamy, Maria Pinto-Albuquerque, and Syed Asad Naqvi. 2019. The Good, the Bad and the Ugly: A Study of Security Decisions in a Cyber-Physical Systems Game. *IEEE Transactions on Software Engineering* 45, 5 (2019), 521–536.
- [17] Yolanda Gil and Donovan Artz. 2007. Towards content trust of web resources. *Web Semantics: Science, Services and Agents on the World Wide Web* 5, 4 (2007), 227–239.
- [18] Deborah L Gladstein and Nora P Reilly. 1985. Group decision making under threat: The tycoon game. *Academy of Management Journal* 28, 3 (1985), 613–627.
- [19] Mark Gondree and Zachary NJ Peterson. 2013. Valuing security by getting [d0x3d!]: Experiences with a network security board game. In *Presented as part of the 6th Workshop on Cyber Security Experimentation and Test*. USENIX, 8.
- [20] Lawrence A Gordon and Martin P Loeb. 2002. The economics of information security investment. *ACM Transactions on Information and System Security (TISSEC)* 5, 4 (2002), 438–457.
- [21] ISO/IEC. 2013. *ISO/IEC 27001*. Technical Report. ISO/IEC. "<https://www.iso.org/obp/ui/#iso:std:iso-iec:27001:ed-2:v1:en>"
- [22] Information Technology Laboratory (ITIL). 2015. *NIST Special Publication 800-53* (v4 ed.). Technical Report.
- [23] Mohammad S Jalali, Michael Siegel, and Stuart Madnick. 2019. Decision-making and biases in cybersecurity capability development: Evidence from a simulation game experiment. *The Journal of Strategic Information Systems* 28, 1 (2019), 66–82.
- [24] Ann C Keller, Chris K Ansell, Arthur L Reingold, Mathilde Bourrier, Mark D Hunter, Sahai Burrowes, and Theresa M MacPhail. [n. d.]. Improving Pandemic Response: A Sensemaking Perspective on the Spring 2009 H1N1 Pandemic. *Risk, Hazards & Crisis in Public Policy* 3, 2 ([n. d.]), 1–37.
- [25] Kari Kelton, Kenneth R Fleischmann, and William A Wallace. 2008. Trust in digital information. *Journal of the American Society for Information Science and Technology* 59, 3 (2008), 363–374.
- [26] James Kendra and Tricia Wachtendorf. 2006. The Waterborne Evacuation of Lower Manhattan on September 11: A Case of Distributed Sensemaking. *Disaster Research Center* (2006).
- [27] Martin G Kocher and Matthias Sutter. 2006. Time is money—Time pressure, incentives, and the quality of decision-making. *Journal of Economic Behavior & Organization* 61, 3 (2006), 375–392.
- [28] Gary McGraw. 1997. Testing for security during development: Why we should scrap penetrate-and-patch. In *Proceedings of COMPASS'97: 12th Annual Conference on Computer Assurance*. IEEE, 117–119.
- [29] Andrew M'manga, Shamal Faily, John McAlaney, and Christopher Williams. 2017. Folk risk analysis: Factors influencing security analysts' interpretation of risk. In *SOUPS 2017*. USENIX Association.
- [30] Tyler Moore, Scott Dynes, and Frederick R Chang. 2015. *Identifying how firms manage cybersecurity investment*. Technical Report. Darwin Deason Institute for Cybersecurity, Southern Methodist University. Available at: <http://blog.smu.edu/research/files/2015/10/SMU-IBM.pdf> (Accessed 2015-12-14).
- [31] John R. Morelock and Zachary Peterson. 2018. Authenticity, Ethicality, and Motivation: A Formal Evaluation of a 10-week Computer Security Alternate Reality Game for CS Undergraduates. In *2018 USENIX Workshop on Advances in Security Education (ASE 18)*. USENIX Association, 11.
- [32] Sai T Moturu and Huan Liu. 2011. Quantifying the trustworthiness of social media content. *Distributed and Parallel Databases* 29, 3 (2011), 239–260.
- [33] National Bureau of Standards, Federal Information Processing Standards Publications (FIPS PUB) 65. 1975. Guideline for automatic data processing risk analysis.
- [34] Jason RC Nurse, Sadie Creese, Michael Goldsmith, and Koen Lamberts. 2011. Trustworthy and effective communication of cybersecurity risks: A review. In *Socio-Technical Aspects in Security and Trust (STAST), 2011 1st Workshop on*. IEEE, 60–68.
- [35] Cabinet Office. 2011. Keeping the Country Running: Natural Hazards and Infrastructure. (2011).
- [36] Alison J Pickard, Pat Gannon-Leary, and Lynne Coventry. 2010. Trust in 'E': Users' trust in information resources in the web environment. In *International Conference on ENTERprise Information Systems*. Springer, 305–314.
- [37] Rex Kelly Rainer Jr, Charles A Snyder, and Houston H Carr. 1991. Risk analysis for information technology. *Journal of Management information systems* 8, 1 (1991), 129–147.

- [38] Ortwin Renn. 1998. The role of risk perception for risk management. *Reliability Engineering & System Safety* 59, 1 (1998), 49–62.
- [39] Ortwin Renn. 2008. *Risk governance: coping with uncertainty in a complex world*. Routledge.
- [40] John R Searle, Ferenc Kiefer, Manfred Bierwisch, et al. 1980. *Speech act theory and pragmatics*. Vol. 10. Springer.
- [41] Adam Shostack. 2014. Elevation of privilege: Drawing developers into threat modeling. In *2014 {USENIX} Summit on Gaming, Games, and Gamification in Security Education (3GSE 14)*.
- [42] Paul Slovic. 1987. Perception of risk. *Science* 236, 4799 (1987), 280–285.
- [43] Joseph St. Germain and Gershon Tenenbaum. 2011. Decision-making and thought processes among poker players. *High Ability Studies* 22, 1 (2011), 3–17.
- [44] William H Starbuck and Frances J Milliken. 1988. Executives’ Perceptual Filters: What they notice and how they make sense. In *The executive effect: Concepts and methods for studying top managers*, Frances J Milliken and William H Starbuck (Eds.). 33–65.
- [45] Rock Stevens, Daniel Votipka, Elissa M. Redmiles, Colin Ahern, Patrick Sweeney, and Michelle L. Mazurek. 2018. The Battle for New York: A Case Study of Applied Digital Threat Modeling at the Enterprise Level. In *27th USENIX Security Symposium (USENIX Security 18)*. USENIX Association, 621–637.
- [46] Detmar W Straub and Richard J Welke. 1998. Coping with systems risk: security planning models for management decision making. *MIS quarterly* (1998), 441–469.
- [47] Lucy Suchman. 2007. *Human-Machine Reconfigurations*. Cambridge University Press.
- [48] Ping An Wang and Easwar Nyshadham. 2011. Knowledge of online security risks and consumer decision making: An experimental study. In *System Sciences (HICSS), 2011 44th Hawaii International Conference on*. IEEE, 1–10.
- [49] Karl E Weick. 1993. The collapse of sensemaking in organizations: The Mann Gulch disaster. *Administrative Science Quarterly* (1993), 628–652.
- [50] Karl E Weick. 1995. *Sensemaking in organizations*. Vol. 3. Sage.
- [51] Darius Wiles and Dave Dugal. 2015. *Common Vulnerability Scoring System v3.0: Specification Document*. Technical Report. FIRST.
- [52] Yan Xu, Evan Barba, Iulian Radu, Maribeth Gandy, and Blair MacIntyre. 2011. Chores Are Fun: Understanding Social Play in Board Games for Digital Tabletop Game Design.. In *DiGRA Conference*.