Are Information Security Professionals Expected Value Maximisers?: An Experiment and Survey-based Test

Konstantinos Mersinas[1], Bjoern Hartig[2], Keith M. Martin[1] and Andrew Seltzer[2, 3]

[1]Information Security Group, Royal Holloway, University of London, Egham, Surrey TW20 OEX, UK

[2]Department of Economics, Royal Holloway, University of London, Egham, Surrey TW20 OEX, UK

[3]Institute for the Study of Labor (IZA), Bonn, Germany

Abstract

Information security professionals have to assess risk in order to make investment decisions on security measures. To investigate whether professionals make such decisions optimally, we conduct an online experiment and survey measuring risk attitudes of security professionals. Participants were asked to state their willingness-to-pay to avoid a series of losses-only lotteries and to make choices between such lotteries. We examine their behaviour in these lotteries and conclude that security professionals do not minimize expected losses. Our findings suggest that security professionals are risk and ambiguity averse and are susceptible to framing effects. We contrast their behaviour to that of a random sample of students. We find that the preferences of security professionals are measurably different from those students in several respects. Finally, we devise a mechanism to elicit professionals' preferences between security and operability. We find that the nature of professionals' employment influences their security versus operability preferences. These factors are usually overlooked in risk assessment methodologies.

1 Introduction

Spending on protective measures for information security is an important issue for most organizations. One of the biggest challenges stems from the fact that specifying the optimal level of information security investment is not an easy task. Reports show that an ever increasing amount of resources is invested on defensive security measures. Nevertheless, there are indications that the cost of security breaches either remains at high levels [51] or has been growing over time [30, 46]. Despite the fact that the budget for security investment is increasing, insufficient expenditure on information security is considered to be one of the main obstacles that security professionals face [30].

One factor limiting security expenditure is a lack of consensus about the optimal approach to decision-making [56]. Professionals are encouraged to choose their own appropriate risk analysis and assessment methods [18, 41] to match the needs of their organisations. Among the approaches used are cost-benefit analysis [35] and risk-management approaches [39]. Even within these broad approaches, there is no single accepted aim; individuals may optimise over Net Present Value, Internal Rate of Return, Return of Investment, and Return of Security Investment [12, 28, 36, 47, 57].

Regardless of approach, existing quantitative risk assessment methodologies are subject to three significant limitations [28]: they are based on many approximations (e.g. unknown risks), these approximations are often biased by perceptions of risk, and the involved calculations can be distorted by the decision-maker's personal biases. These limitations imply that any assessment of appropriate strategies will rely heavily on individuals' judgments. This is not inherently problematic; judgment is unavoidable and is considered necessary for successful risk assessment, as has been recognized in protocol ISO 27005, which states: 'judgment should be exercised in certain cases for the justification of decisions' [41]. However, the importance of individual judgement implies that the decision makers' beliefs, preferences, and ability to process information may play a large role in the optimality of information security decisions. For example, one of the most important factors in a security environment is risk. Decision makers may differ in their tolerance of risk or their assessment of risk in a given environment. The subjectivity of risk perception and the lack of a standard economic model for deciding

on and justifying security investment in an uncertain environment highlight the importance of the decision-maker's preferences and risk attitude. This has been acknowledged within information security. ISO 27005 states: 'consideration should be taken into account about how risk is perceived by affected parties' [41].

Behavioural economics [19, 44, 48] has identified a set of heuristics used by individuals to make decisions. Although heuristics are easy to implement, they do not always lead to optimal decision-making as measured by expected value maximisation. Sub-optimal decision-making may be due to several behavioural biases that frequently characterise heuristics. These biases include loss aversion (excessive focus on negative outcomes), worst case thinking (excessive focus on the worst possible outcome), ambiguity aversion (preference for explicitly stated probabilities or outcomes), risk aversion and inconsistent attitudes over risk, e.g. focussing on the most salient outcomes rather than on all outcomes, and changing behaviour when observed by others. A clear understanding of these potential behavioural biases and their impact on decision making is a useful tool which can lead to the development of appropriate strategies for mitigating the relevant biases.

The nature of decision-making in information security may amplify the importance of several of these behavioural biases. Security professionals regularly decide the amount of protective investment to avoid unwanted losses. It is possible that, even with accurate data on past security breaches and resulting losses, security professionals will invest at suboptimal levels. If an employer is risk neutral, the optimal level of investment will be that which minimises the sum of investment costs and expected losses. However, even if all relevant information is available, a security professional may deviate from this optimal level of investment. When a threat bears potential catastrophic outcomes, the attention of the professional might be disproportionately focused on the worst-case outcome, and hence she might be willing to overspend in order to be on the safe side, even if the probability of such an event is negligible. Alternatively, she might diminish the urgency of quite probable threats or consider small losses inevitable. Additionally, the professional has to balance the level of protection against operational efficiency.

Researchers have used behavioural theories, such as prospect theory [43, 37], in the context of information security [54, 29] but, to our knowledge, actual behaviour of security professionals has not been studied. Our study contributes to understanding active information security professionals' attitudes towards risk and uncertainty, and compares the behaviour of professionals against that of a student sample.[1] *Ex ante* it might be expected that experienced professionals, who routinely work in an environment with multiple risks, would be better at estimating expected outcomes than the sample of students; however, it is also possible that the heuristics used by professionals contain similar biases to the ones used by students. We test this using an experiment designed to elicit both professionals' and students' attitudes towards risk using neutrally framed lotteries.

We have designed our experimental scenarios focussing on several intrinsic attributes of the information security environment, which has operational losses and defence costs and direct losses, in the spirit of [8]. In particular, we focus on the following distinctive set of features, which are examined in our experimental approach:

1. Loss domain: each security investment decision can be described as a lottery with a best outcome of zero. Thus the scope of the decision-maker is loss prevention.
2. Ambiguity of probabilities and outcomes: security professionals face threats that are not precisely known. Often they do not know either the probability of incurring a loss or the likely size of the loss should it occur.
3. Evaluation by other parties: decision-makers in information security need to justify proposed security investment to others, e.g. to business managers or hierarchical superiors.

We find that professionals typically do a somewhat better job of maximising expected value than the student sample, although they too exhibit systematic behavioural biases.

[1] We consider a sample of students randomly drawn from the database of volunteer subjects in the Laboratory for Decision Making & Economic Research at Royal Holloway, University of London, in order to contrast with behaviour of professionals. These students come from all departments of the university.

Specifically, we find evidence that many security professionals are ambiguity averse and change their preferences depending on how a choice is framed.

At the end of the experiment we ask our subjects several survey questions relating to their professional role and to their willingness to trade off security and operability. There is considerable heterogeneity across professionals in their security/operability preferences associated with their professional roles. Most professionals are considerably biased towards one of the two domains and display loss aversion in their preferred attribute.

The rest of the paper is organised as follows. Section 2 presents our behavioural economics theoretical framework and the economic theory behind the experimental and survey approaches. Section 3 presents our core hypotheses and experiment design. Section 4 presents the survey and experimental findings. Section 5 discusses these findings and concludes.

2 Background and Approach

2.1 Behavioural Economics and Security

The study is motivated by the general question of 'how much security is needed?' Our approach focuses on the human factor, specifically risk perception and attitudes of information security decision-makers. The theoretical framework for our approach is from the field of behavioural economics. The importance of behavioural economics to information security has been highlighted in various papers by Anderson and Moore [6, 7, 9, 10]. Several behavioural economics studies have shown that in a variety of contexts agents systematically deviate from the predictions of expected utility theory [27]. This implies that decision making is rarely done by a rational-agent - 'homo economicus' [19, 44, 48].

The starting point for most previous studies has been von Neumann-Morgenstern expected utility theory [27, 24] – which posits that agents work towards a well-defined objective function (such as minimizing the expected losses due to security breaches). Challenges to this approach have come from both psychology [13, 15] and behavioural economics. Bruce Schneier outlines the effect of heuristics and biases on security decision making, describing risk and uncertainty perception issues [5]. Perhaps the most widely used behavioural economics model is prospect theory, which asserts that people value potential gains and losses relatively to a reference point, that losses loom larger than gains and that probabilities are processed in a distorted way [29]. Kahneman and Tversky [43] proposed that decision-makers will be risk-averse for small-probability losses and large-probability gains, but risk-seeking for small-probability gains and large-probability losses. Other cognitive biases relevant for decision making in an information security setting include excessively favouring the status quo (status quo bias) [1, 2, 3]; excessively discounting future outcomes (present bias) [1, 2, 3]; focussing on the most salient outcomes rather than on all outcomes weighted by their probability of occurring [16]; and focussing excessively on worst case, but low probability outcomes [42]. Previous literature has studied the effect of such biases on security design [32] and timing preferences about security investment [40].

One of the defining characteristics of the information security environment is the difficulty of assessing all possible threats. Our approach to measuring how professionals assess threats is drawn from a long literature in economics. We distinguish between risk, ambiguity, and uncertainty in our set-up. This distinction, originally described by Knight [45], concerns the level of information about the underlying environment. *Risk* is defined by all possible outcomes and associated probabilities of a given lottery being known to the decision maker. Only the actual outcome is unknown. *Ambiguity* is defined by imperfect information about probabilities [26]. *Uncertainty* is defined by either outcomes or probabilities being unmeasurable [45]. It is well established that some decision makers are *risk-averse*, *ambiguity averse*, and *uncertainty averse*; implying 1) they prefer a certain outcome to an uncertain one with the same expected value; 2) they prefer a lottery with precisely defined probabilities to one with a range of possible probabilities; and 3) they prefer known distributions of possible payoffs over unknown

ones. One of the factors that affects the extent of risk, ambiguity, and uncertainty aversion is other-evaluation; i.e. the extent to which choices are evaluated by peers, colleagues, superiors, etc. [23]. An explanation for this behaviour is that the individual chooses the most *a posteriori* justifiable option, even if it was not the *ex ante* optimal option.

2.2 Experimental Elicitation of Risk Attitudes

In order to measure risk, ambiguity, and uncertainty aversion, we examine subjects' choices between different lotteries, i.e. lists of consequences with associated probabilities. Following existing literature [59], we operationalise risk, ambiguity, and uncertainty aversion in our experiment by examining subjects' *willingness-to-pay* (WTP), defined as the maximum amount that the individual is willing to sacrifice in order to avoid an undesirable event. Additionally, we present subjects with a series of simple choices between lotteries, in a similar fashion to [38]. All lotteries contain only negative outcomes (as is the case in a security setting). To make our measures incentive compatible, we provide performance-dependent payment to the participants, based on their decisions in these lotteries.

Our approach is a modification of that used by Holt and Laury [38] and is also similar to the alternative of Moore and Eckel [52]. We use the subjects' decisions between sets of lotteries with different expected losses and with different levels of ambiguity in probabilities and in outcomes to create a new instrument for measuring risk and ambiguity aversion. This design allows both comparisons between-subjects, by comparing choices in the same lottery, and comparisons within-subjects, by comparing choices in different lotteries with the same expected value. The advantage of the WTP approach is that it is relatively easy to specify a finite number of possible states that completely determine a lottery's consequences. A key assumption of the standard economic model of 'rational choice under uncertainty' [11] is that where probabilities are not explicitly given, individuals' can assign subjective probabilities to each state of nature. An important point, underlying the experimental elicitation of attitudes, is that

risk and uncertainty, along with the trivial case of certainty, are a means to reveal the decision-maker's preferences, and also their *belief* in how plausible events are to occur ([17], Ch.15.2.2).

Appendix B.5 shows the instructions given to subjects in the experiment along with screen shots of the experiment.

2.3 Surveys

This study combines an incentivized experiment with non-incentivized survey data. Our survey asks questions on professional experience, such as the role in which professionals are currently employed and the years of experience in information security related tasks.

Experiments and surveys both have advantages and disadvantages, and we combine the two approaches to increase the reliability of our results. The primary advantage of incentivized experiments is that subjects have a stake in their answers, so the data is less likely to be subject to biases associated with people responding differently to hypothetical situations [50]. The primary advantage of surveys is that they can be designed to examine a real world context. However, survey data is not generated by incentivised individuals. Moreover, survey respondents may misunderstand questions, have difficulty recalling information, or be influenced by socially acceptable answers.

Our results are drawn primarily from the experiments, because incentivised decisions are more likely to accurately reflect underlying preferences. However, we also check whether observed behaviour during the experiment is correlated with the answers to the survey questions. Specifically, we examine how professional experience correlates with observed risk attitude throughout the experiment. This follows the results in [25], which shows that risk attitudes in survey data and experiment input were strongly correlated across subjects.

Appendix A.1 shows the survey that subjects answered after completing the experiment.

3 Methodology

3.1 Hypotheses

We analyse the behaviour of security professionals and students in our experiment and survey to test a series of hypotheses motivated by the following commonly observed behaviour patterns:[2]

1. *Risk and ambiguity aversion*: Risk aversion implies that given a lottery with a small probability of loss, an individual is willing to pay more than the expected value of these losses to avoid playing the lottery. Ambiguity aversion implies that for a lottery with the same expected losses, an individual is willing to pay an additional amount above the risk premium to avoid the lottery if, instead of a specified probability or outcome, there is a range of probabilities or outcomes. However, prospect theory implies that for larger probabilities of losses, the same individuals may engage in risk-seeking behaviour [43]. It is possible that security professionals differ systematically from the student population in regards to risk and ambiguity aversion because the nature of their work implies greater exposure to risk and ambiguity. Security professionals face continual threats of losses, which are often not well defined.

2. *Worst-case aversion:* This implies that individuals pay disproportionate attention to the worst possible outcomes [16]. Their WTP to avoid playing a lottery increases in the maximum possible loss, even if the expected value and variance of a lottery is held constant. Worst-case thinking may be particularly common among security professionals, as the field has seen a number of high-profile cases of catastrophic losses due to security breaches in recent years [55]. On the other

[2] Loss aversion, i.e. a disproportionate weighting given to outcomes of less than zero, is another anomaly that has received considerable attention in the behavioural and experimental economics literatures. We do not focus on loss aversion because the information security environment involves losses only.

hand, small losses due to security breaches may be regarded as a normal part of the operating environment and not be worthy of any expenditure.

3. *Other-evaluation*: This implies that when decisions are evaluated by other parties, individuals tend to be more risk averse, ambiguity averse, and worst-case averse. Since evaluators do not observe *ex ante* probabilities, only *ex post* outcomes and thus may blame the decision-maker for bad outcomes even if the decision that led up to it was *ex ante* correct 'a decision maker, […] makes the choice that is perceived to be most justifiable to others.' [23] Other-evaluation may be particularly important in a security context, as security decision-makers normally have to justify their investment proposals to business managers, chief officers, the board of directors, etc.

We examine these behaviour patterns for both professionals and students in our experiments. In the case of security professionals, we also explore a fourth aspect of decision making in the survey part:

4. *Security and Operability*: We expect that security professionals will tend to value security more than operability. In other words, when balancing the costs of implementing security controls against the resulting loss of efficiency of business operation, security professionals will select a trade-off position that prioritizes security ahead of operability.

3.2 Experimental Procedure

We conducted the experiment with two different samples. The sample of information security professionals was drawn from current and previous students of the distance learning MSc in Information Security at Royal Holloway, University of London (RHUL) and consisted of 59 individuals (53 male) with an average age of 38. This group consists of security professionals who work in the industry and were undertaking the distance learning master's program on a part-time basis. The mean industry experience of this group is 8.64 years. The student sample was drawn from individuals registered in the database of the Laboratory for Decision Making and Economic Research at RHUL. This

group consists of 58 active full-time students (24 male) from all departments of the university with an average age of 21.4. [3]

Our experiment consists of several lotteries designed to test our hypotheses. The lotteries were framed neutrally for two reasons. First, we are trying to measure the underlying preferences of security professionals, not their interpretation of professional standards regarding threats. Secondly, the neutral framing means that student subjects are not being asked to make decisions on matters they have never previously experienced. Thus the student and professional samples can be considered directly comparable. All lotteries in the experiment require 'one-off' decisions, with no feedback given after a choice is selected.

One set of lotteries elicits risk and ambiguity attitudes across three levels of expected losses and three levels of probabilities. Specifically, subjects are asked to choose between lotteries where ambiguity of both probability and loss are changed one at a time, or simultaneously.[4] This approach enables us to compare WTP between-subjects and also within-subjects across different types of risky and ambiguous decisions.

In another set, the lotteries differ from each other in terms of worst-outcome, expected value and variance. These lotteries allow us to examine whether subjects employ simple decision rules (heuristics) to choose between the complex lotteries. Additionally, we elicit both WTP and binary choices for a subset of these lotteries, allowing us to check whether our subjects' preferences are consistent across different framings.

Furthermore, it is possible that security professionals have a tendency to overemphasise security issues at the expense of operational issues. To examine this question, we ask subjects to choose between security and operability in a realistic scenario. Our approach follows [20]. Operability is framed as the time needed for task completion. Operability and security are both framed explicitly in monetary terms. To exclude other factors, the

[3] Three subjects in the student sample were excluded from the analysis because they stated that they were related to information security.

[4] Participants were informed that they would receive a fixed participation payment and an additional potentially larger amount depending on their decisions in the experiment. In particular, one of the lottery comparisons of Appendix A.2 was randomly chosen for each participant, and their preferred lottery was 'played' by a pseudorandom probability generator. The outcome was mapped to a maximum performance gain of 10 USD and was sent along with the participation payment to individuals, in the form of an Amazon gift certificate.

scenario describes an information system of moderate-impact to confidentiality, integrity, and availability [58].

Finally, subjects fill out a short questionnaire about their personal attitudes and demographics. We use this data to examine correlations with behaviour in the main experiment.

3.3 Experiment design

3.3.1 Risk and ambiguity aversion

We test risk and ambiguity preferences using 12 neutrally framed lotteries, divided into three groups of four (see Table 9 in Appendix A.2). Lotteries within a group have identical expected value, but different degrees of ambiguity. Subjects are asked to state their maximum WTP in order to avoid playing each lottery. For example, the four lotteries in group A ($H_1 1$ to $H_1 4$) all have an expected value of $\mu = -2.5$ and contain the following text:

'What is the maximum amount that you are willing to pay in order to avoid playing a lottery in which there is ..:

i. *..a 5% probability of losing $50 and losing nothing otherwise?'*
ii. *..a probability between 0% and 10% of losing $50 and losing nothing otherwise?'*
iii. *..a 5% probability of losing between $20 and $80 and losing nothing otherwise?'*
iv. *..a probability between 0% and 10% of losing between $20 and $80 and losing nothing otherwise?'*

The lotteries in group B ($H_1 5$ to $H_1 8$) and group C ($H_1 9$ to $H_1 12$) contain the same potential outcomes, but probabilities of 15% (0%-30%) and 50% (35%-65%), respectively. Hence the ambiguous lotteries are all designed such that if there was a uniform distribution of outcomes and probabilities over the range of ambiguity, the expected value of losses would be the same as in the risky lottery. Appendix A.2. contains the full set of lotteries. In the following, we refer to lotteries of type i. as *risky*,

lotteries of type ii. as *probability ambiguous*, lotteries of type iii. as *outcome ambiguous*, and lotteries of type iv. as *fully ambiguous*.

Subjects had to give their WTP for all 12 lotteries, but the order in which is lotteries were presented was counterbalanced to control for potential order effects. That means that some subjects saw the risky lotteries first and some saw the fully ambiguous lotteries first.

3.3.2 Worst-case thinking and other heuristics

To test the *worst-case aversion* hypothesis, we ask subjects to choose between pairs of lotteries, with each lottery containing five outcomes. The probabilities are identical across all lotteries, but outcomes were different. We chose the outcomes in such a way that the expected value is identical in some pairs and different in others. For example, lottery L_2 (L_3) contains the following outcomes:

- 15% probability of losing nothing (nothing)
- 30% probability of losing 166.66 (183.33)
- 30% probability losing 300 (300)
- 20% probability of losing 450 (450)
- 5% probability of losing 900 (800)

While both lotteries L_2 and L_3 have the same expected value (-275), the highest loss in lottery L_2 (900) is greater than in lottery L_3 (800). In other words, lottery L_3 contains the 'worse worst case'.

Instead of worst-case thinking, subjects may also use other heuristics or simple decision rules to decide between lotteries. For example, subject may put a lot of weight on the best possible outcome ('best-case thinking'). Or they may pairwise compare states across lotteries and prefer the lottery which 'wins' in more states. Finally, subjects may also prefer lotteries with less variance, which would constitute a form of risk aversion. In order to test whether subjects use any of these heuristics, we compare the majority

choice in our samples with the predictions of each heuristic. If any heuristic is consistent with all or at least most of the majority choices, it would provide evidence that subjects indeed rely on simple decision rules.

In total we have eight different lotteries which are used in five pairwise comparisons (two lotteries are used twice); three with an expected value of -275 and five with expected values ranging from -86.25 to -86.75. Appendix A.3 contains further details.

In addition to the pairwise choices, we also elicit subjects' WTP for three of the eight lotteries (L_5, L_6 and L_7, see Appendix A.4). Since these three lotteries are also used in two pairwise choices, it allows us to check whether our subjects' preferences are influenced by the type of decision. Such inconsistencies would violate rational choice theory since rational preferences should be unaffected by the way in which they are elicited (choice or WTP). The three WTP questions are separated from the pairwise choices by a different unrelated set of questions in order to disguise the similarities of the decisions, and both types of questions were counterbalanced.

All eight lotteries are designed to approximate power-law distributions. Such distributions simulate the occurrence of rare events that are observed in various physical and social phenomena, such as earthquakes, war deaths, web hits, and citations [53] and have also been observed in security issues, such as identity theft and malware spreading [33, 49]. In the general form of a power-law distribution, probability p is specified as a function of outcome $(x) = \frac{\kappa}{(-x)^{\alpha}}$, where α is the distribution exponent and κ a constant. For the purposes of our experiment, we set $\alpha = 1.1$, constant $\kappa = 20$ and $x \in [-1000, 0]$.

3.3.3 Other-evaluation and behaviour

We examine other-evaluation using a between-subjects design in which subjects are assigned to either a *control group*, which is presented with the standard version of the experiment, and a *treatment group*, which is initially informed that all choices made in the experiment would be 'further viewed' and would 'go through an additional evaluation process'. Participants are informed that the evaluators would have the same

information as themselves [21]. See Appendix B.5 of the supplementary data for the original instructions.

Ultimately any test of other-evaluation in an experiment such as this is going to be fairly weak for two reasons. First, the experiment itself has fairly low stakes, so any evaluation done within the experiment will not have much impact. This alone may not prevent other-evaluation from impacting subjects' behaviour [4]. Additionally, however, since our experiment was conducted on-line, we could not give any public feedback, limiting the perceived social impact of the evaluation.

3.3.4 Relative importance of security and operations

This part of the study consists of two sets of questions given to the professional sample only. Subjects are asked scenario-based questions in which they have to choose between two measures with different impact on the security level and the operability of a system. Both security and operability have equal monetary values assigned to them. The specific questions asked are:

> *'Imagine the following scenario: You are managing an Information System that has moderate-impact on the confidentiality, availability and integrity of information records kept by your organisation.*
> *The total worth of the system under protection is evaluated at $10,000.*
> *Full operability of the system allows the business to gain a profit of $10,000.*
> *Two new mechanisms A and B with the same cost are proposed for the system.*
> *Which one of the following mechanisms do you prefer?' (Table 1)*

Table 1: Initial question of Scenario 1: 'Which one of the following do you prefer?

Mechanism A	Mechanism B
Enhances Security of the system by 10%	Enhances Operability of the system by 10%

Subsequent questions are formed dynamically depending on previous answers. In the next question the value of the preferred measure is marginally decreased while the value of the other measure remains constant. This is repeated until the subject crosses over from choosing one measure to the other, so that a switching point between security and operability is specified.

The second set of questions elicits a measure for whether losses of the attribute preferred in Scenario 1 (security or operability) are treated differently than gains.[5] Subjects are asked to choose between three options (Table 2):

Table 2: Scenario 2 template question

Choice A	Choice B	Choice C
Remains at the current system state	Reduces Security by *x*% Enhances Operability by *y*%	Indifferent between A and B

Values *x* and *y* of choice B are taken from the switching point that is computed from Scenario 1. If a subject selects choice B or C, this stage of the experiment ends. If the subject chooses A, then the question is repeated, except if operability (security) has been preferred in the previous scenario, the security (operability) reduction is lowered by one percent. To illustrate, consider the following example: In Scenario 1 a subject is indifferent between a 5% security enhancement and a 10% operability enhancement. Subsequently, in Scenario 2, choice B gives a 5% reduction in security and a 10% enhancement in operability. If choice A is selected (i.e. choice B is considered worse than the status quo), the reduction in security in choice B is decreased to 4%, and so on.

The difference between the values of Scenario 1 and 2 (if any) constitutes our measure of loss aversion on the preferred attribute (security or operability). If losses and gains of equal magnitude are weighed equally, subjects will always select choice C. If, however, a loss looms larger than an equivalent gain, subjects will prefer choice A. Further details on the choice of Scenarios 1 and 2 can be found online in the supplementary data (Appendix B.1).

[5] Other questions unrelated to security and operability were asked between scenarios 1 and 2 to disguise the relationship between the two scenarios.

4 Analysis and Findings

This section presents our findings for the main hypotheses outlined in the previous section. Following standard experimental economics procedures, the experiment is counterbalanced to control for potential order effects, data has been checked for validity and cleaned accordingly[6] and outliers have been shown to be non-influential (see supplementary data – Appendices B.2, B.3 and B.4) for the relevant tests.

4.1 Risk and Ambiguity Aversion

Table 3 lists subjects' average WTP to avoid playing a lottery for the 12 lotteries $H_1$1 to $H_1$12 for both professional and student sample. Figure 1 depicts the difference between the expected loss of a lottery and the subjects' average WTP. Positive (negative) values indicate that subjects are willing to pay more (less) than the lottery's expected loss.

> ***Finding 1****: Both professionals and students are risk averse for small probability losses. Professionals become risk seeking for high probability losses.*

Both security professionals and students are willing to pay significantly more (two-sided t-test, $p < 0.01$) than the expected value of the two *risky* lotteries $H_1$1 and $H_1$5, which have a probability of a loss of 5% and 15%, respectively. However, security professionals pay (weakly) significantly less ($p < 0.1$) than the expected value when the probability of a loss is 50% in $H_1$9. Students do not significantly deviate from the expected value in this lottery. This behaviour is consistent with prospect theory, which

[6] For example, we have removed subjects who were willing to pay for than $50 to avoid playing a lottery with a maximum loss of $50.

predicts risk aversion in small probability losses, but more risk seeking behaviour when probabilities of losses are large.

Figure 1: Difference between subjects' WTP and expected value

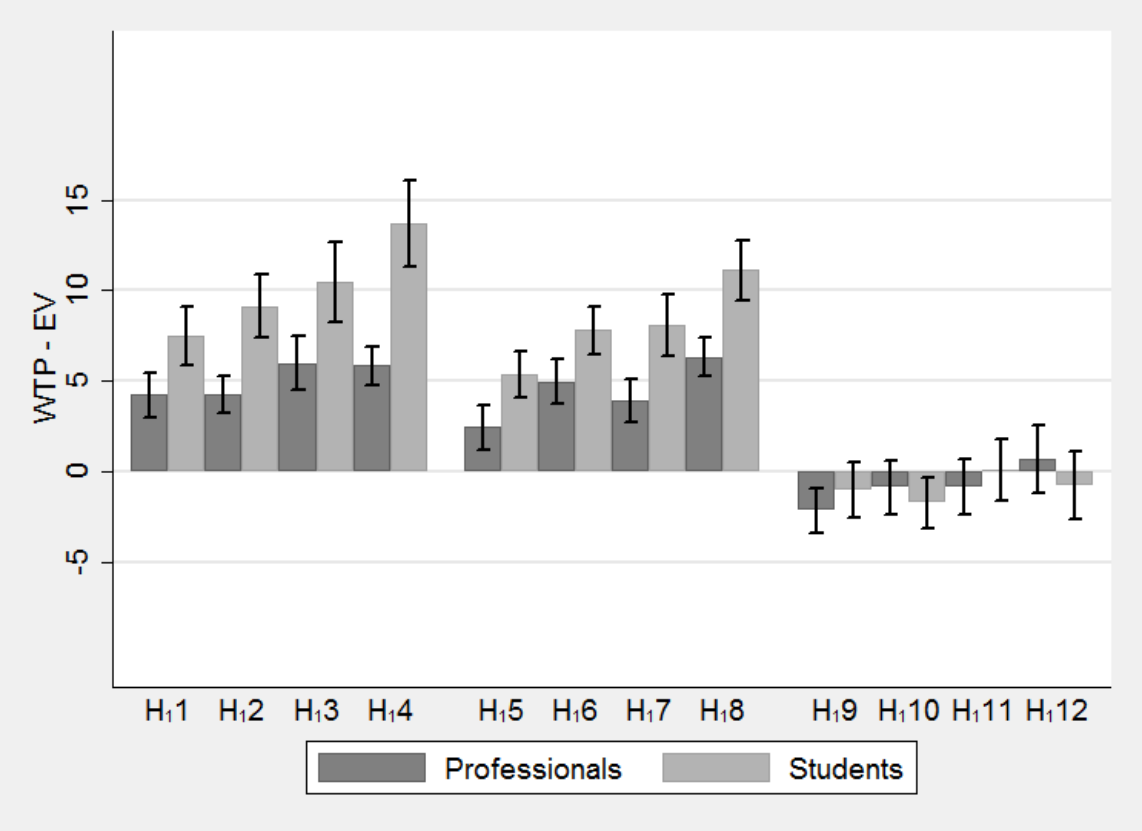

Note: Bars represent participants' mean WTP minus the EV of each of the 12 lotteries.

***Finding 2**: Professionals reveal ambiguity aversion in all of their choices; ambiguity aversion is not consistently observed in the student sample.*

To measure ambiguity aversion, we compare our subjects' WTP in the *risky* lotteries ($H_1$1, $H_1$5, $H_1$9) with their WTP in the three corresponding *ambiguous* lotteries ($H_1$2-4, $H_1$6-8, $H_1$10-12). If ambiguous outcomes and probabilities were uniformly distributed, these lotteries would have the same expected value as the risky lottery. Therefore, if subjects are willing to pay more to avoid an ambiguous lottery than the corresponding risky lottery, we consider them to be ambiguity averse.

We find that professionals are significantly more likely to pay more to avoid a *fully ambiguous* lottery ($H_1$4, $H_1$8, $H_1$12) than to avoid a *risky* lottery in all three groups (sign-rank test, A, B: $p < 0.01$, C: $p < 0.05$).

The effect of ambiguity is less strong when only one of the elements (probability or outcome) is ambiguous ($H_1$2-3, $H_1$6-7, $H_1$10-11). Although the general tendency observed in these decisions is consistent with ambiguity aversion, only three of the six comparisons with the risky lotteries are statistically significant (see Table 3).

Students are even more ambiguity averse when loss probabilities are low (A, B: $p < 0.01$)[7], but not when the loss probability is high, with the (non-significant) average effect going in the opposite direction.

Neither the professional nor the student sample show significant differences between *probability ambiguous* and *outcome ambiguous* lotteries.

> ***Finding 3****: Professionals seem to deviate less from expected value maximization than the student sample when probabilities are low and lotteries are ambiguous.*

Generally, professionals' WTP is closer to the expected value than students' WTP. This is the case in 10 out of the 12 lotteries (Table 3), although the differences are very small in group C ($H_1$9 – $H_1$12). The between-sample comparison is statistically significant in five of eight lotteries of groups A and B (see Table 3). All these five lotteries are ambiguous, suggesting that professionals are better in dealing with ambiguity when loss probabilities are low. This may be due to their work experience, where they are naturally confronted with ambiguity and loss-probability events.

[7] In the student sample, even a two-sided paired t-test against the WTP in the risky lottery is highly significant in group A and B, which is not the case in the professional sample.

Table 3: Subjects' WTP to avoid lotteries $H_1 1 - H_1 12$

	Professionals (N=59)			**Students (N=58)**	
	EV	WTP (SE)	⇔	EV	WTP (SE)
$\boldsymbol{H_1}$ **1**	-2.5	6.71^{***} (1.22)		-2.5	9.97^{***} (1.64)
$\boldsymbol{H_1}$ **2**	-2.5	6.76 (1.01)	$^{\phi\phi}$	-2.5	$11.62^{\#}$ (1.76)
$\boldsymbol{H_1}$ **3**	-2.5	$8.47^{\#\#}$ (1.52)	$^{\phi}$	-2.5	$12.95^{\#\#\#}$ (2.20)
$\boldsymbol{H_1}$ **4**	-2.5	$8.34^{\#\#\#}$ (1.05)	$^{\phi\phi\phi}$	-2.5	$16.19^{\#\#\#}$ (2.35)
$\boldsymbol{H_1}$ **5**	-7.5	9.92^{*} (1.21)		-7.5	12.81^{***} (1.27)
$\boldsymbol{H_1}$ **6**	-7.5	$12.46^{\#\#\#}$ (1.24)		-7.5	$15.29^{\#\#\#}$ (1.32)
$\boldsymbol{H_1}$ **7**	-7.5	$11.41^{\#\#}$ (1.18)	$^{\phi\phi}$	-7.5	$15.57^{\#\#}$ (1.74)
$\boldsymbol{H_1}$ **8**	-7.5	$13.81^{\#\#\#}$ (1.08)	$^{\phi\phi}$	-7.5	$18.59^{\#\#\#}$ (1.68)
$\boldsymbol{H_1}$ **9**	-25	22.81^{*} (1.24)		-25	23.93 (1.52)
$\boldsymbol{H_1}$ **10**	-25	$24.10^{\#}$ (1.49)		-25	23.24 (1.38)
$\boldsymbol{H_1}$ **11**	-25	24.15 (1.54)		-25	25.10 (1.70)
$\boldsymbol{H_1}$ **12**	-25	$25.64^{\#\#}$ (1.85)		-25	24.22 (1.86)

* $p \leq 0.1$, ** $p \leq 0.05$, *** $p \leq 0.01$; indicates significant difference of WTP and EV in two-sided t-test for given lottery.
$^{\#}$ $p \leq 0.1$, $^{\#\#}$ $p \leq 0.05$, $^{\#\#\#}$ $p \leq 0.01$; indicates significant difference of WTP in given lottery and corresponding risky lottery ($H1_1$, $H1_5$ or $H1_9$)in two-sided sign-rank.
$^{\phi}$ $p \leq 0.1$, $^{\phi\phi}$ $p \leq 0.05$, $^{\phi\phi\phi}$ $p \leq 0.01$; in two-sided t-test comparing WTP between samples.

4.2 Worst-case thinking and other heuristics

This section is divided into the analysis of two parts: pairwise lottery comparisons and lottery comparisons against stated WTP.

4.2.1 Lottery Comparisons and findings on potential heuristics

***Finding 4**: There is no single decision rule that explains a large majority of choices.*

We ask subjects to make five pairwise choices between eight lotteries varying in expected value, variance, number of dominant states, best (least bad) and worst outcome (see Appendix A.3). Table 4 summarizes the choices and Table 5 shows the extent to which the majority choice is consistent with different simple decision rules based on these five criteria (expected value, variance, number of dominant states, best and worst outcome).

We do not find much evidence for behaviour systematically favoring one of the decision rules. Although choices seem to be most consistent with expected value maximization or state-by-state comparison, the results do not sufficiently tilt into any one direction to draw meaningful conclusions. In fact, only one of the choice patterns is statistically significantly different from 50-50.

Table 4: Professionals' and students' choices between pairs of lotteries and across sample comparisons.

Comparison variable	Choice pair	Professionals: (N=59)	Students (N=58)	Comparison Professionals vs Students[#]
H_2 1	L_5 VS L_6	L_5: 35 (59%) L_6: 24 (41%)	L_5: 29 (50%) L_6: 29 (50%)	Difference 9%
H_2 2	L_6 VS L_7	L_5: 31 (53%) L_6: 28 (47%)	L_5: 35 (60%) L_6: 23 (40%)	Difference 7%
H_2 3	L_4 VS L_2	L_5: 21 (36%)** L_6: 38 (64%)**	L_5: 28 (48%) L_6: 30 (52%)	Difference 12%
H_2 4	L_2 VS L_3	L_5: 36 (61%) L_6: 23 (39%)	L_5: 35 (60%) L_6: 23 (40%)	Difference 1%
H_2 5	L_1 VS L_8	L_5: 32 (54%) L_6: 27 (46%)	L_5: 30 (52%) L_6: 28 (48%)	Difference 2%

* $p \leq 0.1$, ** $p \leq 0.05$, *** $p \leq 0.01$; indicates the observed share is significantly different from 50% (Binomial probability test with p = 50%).

[#] For each of the five pairs, the distributions among professional and among students were tested for significant differences using a Fisher-exact test. None of comparisons yielded statistically significant differences.

Table 5: Professionals' and students' choices between pairs of lotteries and consistency with heuristics

Lottery pair	Expected Value	Variance	Worst outcome	Best outcome	# of dominant states
	Professionals				
L_5 VS L_6 (59%, 41%)	–	×	✓	✓	–
L_6 VS L_7 (53%, 47%)	✓	✓	×	×	✓
L_4 VS L_2 (36%, 64%)	–	×	×	✓	–
L_2 VS L_3 (61%, 39%)	–	×	×	✓	–
L_1 VS L_8 (54%, 46%)	✓	✓	×	✓	✓
	Students				
L_5 VS L_6 (50%, 50%)	–	–	–	–	–
L_6 VS L_7 (60%, 40%)	✓	✓	×	×	✓
L_4 VS L_2 (48%, 52%)	–	×	×	✓	–
L_2 VS L_3 (60%, 40%)	–	×	×	✓	–
L_1 VS L_8 (52%, 48%)	✓	✓	×	✓	✓

✓: indicates that the majority choice is consistent with the heuristic's prediction (this does not indicate statistical significance!).
×: indicates that the majority choice is inconsistent with the heuristic's prediction (this does not indicate statistical significance!).
-: indicates that the heuristic does not make a unique prediction for the lottery or that there is no majority choice.

4.2.2 Consistency across types of decision

We elicit subjects' WTP for L_5, L_6, and L_7, which are also used in two pairwise choices (see Appendix A.4). This allows us to check our subjects' decisions for internal consistency. The results are summarized in Table 5.

> ***Finding 5****: Security professionals exhibit preference inconsistencies between willingness-to-pay and choice decisions. The level of inconsistency is similar to the student sample.*

A considerable share of our professional subjects exhibit preference reversal in both comparisons (52.5% and 33.9%), which is similar to the share of preference reversals observed in the student sample (43.1% and 41.4%). Fisher-exact tests do not reveal any significant differences between the two samples. This suggests that professionals are as susceptible to be (irrationally) affected by the way in which a decision is presented to them as students.

Table 6: Lottery comparisons and willingness-to-pay inconsistencies

Professional sample (N=59)

		Choice	
		$L_5 \succ L_6$	$L_6 \succ L_5$
WTP	$L_5 \succ L_6$	16*	6
	$L_5 \sim L_6$	2	3
	$L_6 \succ L_5$	17	15*

		Choice	
		$L_6 \succ L_7$	$L_7 \succ L_6$
WTP	$L_6 \succ L_7$	11*	17
	$L_6 \sim L_7$	3	2
	$L_7 \succ L_6$	17	9*

Student Sample (N=58)

		Choice	
		$L_5 \succ L_6$	$L_6 \succ L_5$
WTP	$L_5 \succ L_6$	9*	10
	$L_5 \sim L_6$	0	3
	$L_6 \succ L_5$	20	16*

		Choice	
		$L_6 \succ L_7$	$L_7 \succ L_6$
WTP	$L_6 \succ L_7$	20*	16
	$L_6 \sim L_7$	2	3
	$L_7 \succ L_6$	13	4*

* Inconsistent choice

4.3 Other-evaluation Ambiguity Aversion

> ***Finding 6****: There is no evidence that subjects change their risk behaviour when they are informed that they will be evaluated by other parties.*

We observe no significant differences between the control and the treatment groups in either lottery comparisons or WTP questions, although in other experiments [4] other evaluation has been shown to have a large effect on the behaviour of subjects. [8] A possible explanation for the absence of an effect here is that our treatment manipulation was unsuccessful. It might be impossible to create a heightened sense of 'being evaluated by another party' in an online study, especially when participants already know that their responses will be subject to 'evaluation' for statistical analysis. It is possible that individuals would be much more concerned with other-evaluation in a different environment, and we believe that this is an important area for future research.

4.4 Security-Operability trade-off

We ask our professional subjects to choose between two mechanisms of equal value, one enhancing security and the other enhancing operability. Then we decrease the value of the preferred mechanism until a subject becomes indifferent.

> ***Finding 7****: These preferences of security professionals are to a great extent dependent on their job role.*

The majority of professionals (58%) prefer a mechanism enhancing the operability of a system over one enhancing its security. However, the actual information security roles of professionals noticeably influences their preferences. Generally, compliance and risk professionals, as well as senior executives, tend to be security-oriented. On the other hand, professionals with managerial roles tend to prefer operability, while IT professionals as a group did not lean strongly in any direction (Table 6).

[8] We use both parametric (e.g. t-test) and non-parametric tests (e.g. Mann-Whitney-U-test) to compare WTP and choices between the treatment and the control group. We find no significant differences anywhere, which is why we do not report the test results explicitly.

Table 7: Security VS Operability preference across job titles

	Job Title					
	Senior executive role[1]	**Managerial role[2]**	**IT & Security role[3]**	**Compliance, Risk or Privacy role[4]**	**Other**	**Total (N=56)[5]**
Mechanism A Enhances Security of the system by 10%	4	3	8	8	2	25
Mechanism B Enhances Operability of the system by 10%	2	13	8	3	5	31

[1] e.g. CEO, CIO, CISO, CSO etc.
[2] e.g. Project Manager, IT Director, Security Manager etc.
[3] e.g. Security Officer, System Administrator, Cyber Security Information Analyst etc.
[4] e.g. Governance, Risk & Compliance Consultant, Information Security Consultant, Auditor etc.
[5] Three subjects did not answer this question.

Finding 8*: The preferences for either security or operability are non-negligible.*

The majority of subjects do not switch immediately back to the other mechanism once the preferred mechanism is made less attractive (see Figures 2 and 3). This suggests that subjects do in fact possess strong preferences for either security or operability. On average, subjects switch after their preferred mechanism has been reduced from 10% to 5.9% for those preferring security, and 5.5% for those preferring operability.

Finding 9*: Professionals tend to weigh losses in their preferred attribute more strongly than gains.*

The second measurement is the relative loss aversion with respect to the preferred attribute, as described in Section 3.3.4. We find that more than half of all professional subjects for whom loss aversion can be meaningfully elicited[9] (32 of 50) display some degree of loss aversion, irrespective of whether they prefer security (15 of 22) or operability (17 of 28). On average, subjects become indifferent once the reduction in

[9] Loss aversion can only be elicited if a trade-off between both attributes exists which makes the subject indifferent.

their preferred attributed is lowered by more than 2% (SEC: 2.2%, OPS: 2.1%), i.e. they value a loss of X% in their preferred attribute about as much as a X+2% gain.

We do not find statistically significant differences between the subjects who prefer security and those who prefer operability, neither with respect to the share of loss averse subjects (two-sided Fisher exact, p = 0.40) nor with respect to the strength of their loss aversion (Mann-Whitney-U-test, p = 0.79). See also Figures 4 and 5.

Figure 2: Switching point when security is preferred (*Sec: x%, Ops:10%*)

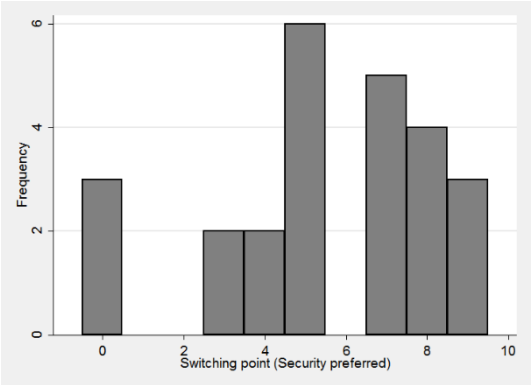

Figure 3: Switching point when security is preferred (*Sec: 10%, Ops:x%*)

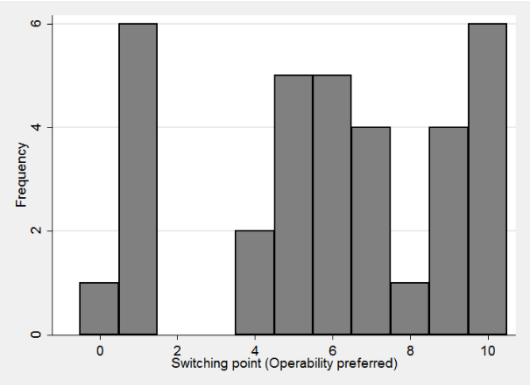

Figure 4: Switching point loss aversion when security is preferred (*Sec: -x+y%, Ops:10%*)

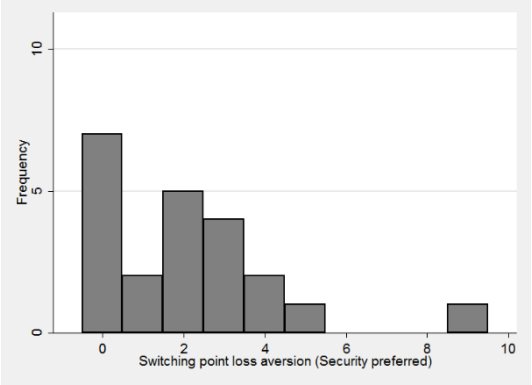

Figure 5: Switching point loss aversion when security is preferred (*Sec: 10%, Ops:-x+y%*)

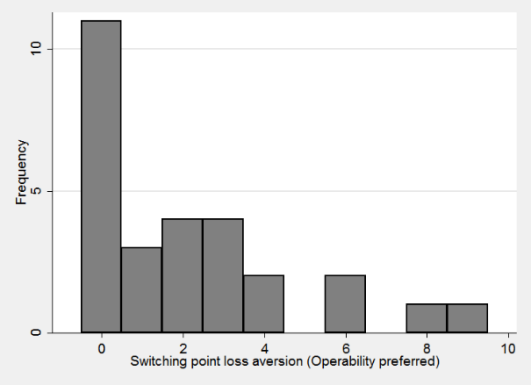

4.5 Survey Analysis

After the main experiment, we ask subjects to fill out a questionnaire about attitudes, opinions and demographics. We now present a couple of notable correlations from these questions. While we do not consider these results among the primary contributions of our study, they are nevertheless interesting and might motivate further research in the future.

> ***Finding 10****: Security professional reveal different risk attitudes to the ones they self-report.*

Subjects rate their willingness to take risks in general on a scale from 0 (not at all prepared to take risks) to 10 (fully prepared to take risks). The student sample confirms the previous finding by [25] that self-reported risk attitude and behaviour is correlated. The professional sample, however, shows the opposite result. Higher self-reported willingness to take risks tends to be correlated with higher WTP to avoid playing lotteries, i.e. higher risk-aversion. See Table 7 for the complete set of correlation coefficients and Figures 6 for a visualization of the sample differences for an exemplary lottery.

Figure 6: Interaction of *Professional* or *Student* and H_19 with *General Risk* as moderator.
'High' denotes a risk seeking individual (self-reported) and
'Low' denotes a risk averse individual (self-reported)

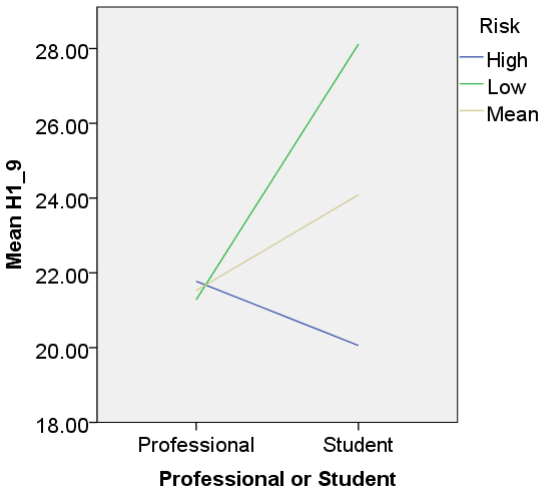

> ***Finding 11****: The number of family depends of participants influences risk attitude of professionals and students differently.*

Table 8: Spearman's correlation coefficients between self-reported risk attitude and WTP

	Professionals	**Students**
	Rho (p)	Rho (p)
H_1 1	.117 .378	-.030 .823
H_1 2	.131 .324	-.080 .550
H_1 3	.227* .085	-.086 .520
H_1 4	.303** .020	-.113 .400
H_1 5	.213 .291	-.080 .550
H_1 6	.291** .025	-.088 .512
H_1 7	.279** .032	-.177 .183
H_1 8	.363*** .004	-.114 .393
H_1 9	.007 .616	-.266** .044
H_1 10	.131 .322	-.252* .057
H_1 11	.005 .972	-.181 .174
H_1 12	.008 .952	-.187 .160

Summary and Discussion

Information security is a field with inherent risk and uncertainty. Organizations and policy makers have sought to reduce the magnitude of these issues; for example, by gathering data on historical security breaches or passing new disclosure laws which increase public knowledge about the distribution of breaches (such as the California Security Breach Information Act [14]). Despite these efforts to collect information, the complexity and uniqueness of information security systems often only allow organizations to approximate *ranges* of probabilities and of damages associated with potential threats and vulnerabilities. Thus, risk management and security investment are, by nature, characterized by ambiguity and uncertainty. This study examines how information security professionals make decisions in such an environment and, specifically, whether security professionals are rational decision-makers who minimize expected losses.

Expected utility theory is the standard normative approach to decision making. Expected utility theory states that for decisions that are made frequently, a rational decision maker should minimize expected losses. However, behavioural economics has repeatedly demonstrated that most individuals systematically deviate from expected utility maximization. In this study, we examine three well-known behavioural anomalies: risk and ambiguity aversion, worst-case aversion, and other-evaluation. We also examine a fourth, industry-specific behavior, namely a preference for security over operability. We examine these behaviours using an experiment and survey that elicits preferences using simple, neutrally-framed lotteries. We compare decision-making of professionals in the experiment to a sample of university students.

Across a variety of lotteries, information security professionals consistently indicated a willingness to pay to avoid negative outcomes that was closer to the expected losses than did the sample of students. This suggests that they are better able to accurately assess risks. One interpretation of this is that their ability to assess risks and minimize the consequence of threats has been shaped by the constant exposure to risk inherent to the security environment.

Despite their greater ability to assess risk, our findings suggest that security professionals still have distinctive behavioural characteristics that deviate from expected utility theory. In common with the student sample, and with a number of other studies, the observed behaviour

of professionals follows the pattern of risk attitudes described by Kahneman and Tversky [43]. Security professionals exhibited significant risk aversion when faced with low possibilities of loss or small losses. However, their actions switched from being risk averse to being risk seeking when faced with large probabilities of losses or large losses. This behaviour is not consistent with expected utility theory, and has particularly important implications in an information security environment where small losses may be inevitable and there exists an ever-present, but low, probability of catastrophic losses. The combination of risk-averse behaviour for small-losses and risk-seeking behavior for large losses could result in over-investment in simple preventive measures for common information security threats (e.g. malware, viruses); but under-investment in measures against potentially catastrophic breaches.

Information security professionals also showed considerable ambiguity aversion in the experiment. Their WTP increased significantly when faced with low probability lotteries which had ambiguous probabilities or outcomes. The extent to which WTP increased did not depend on whether the ambiguity was in the probability that an outcome occurs or whether it was in the cost of the outcome. As with risk, ambiguity is an inherent feature of the information security environment, which is characterized by unknown or imperfectly known threats.

Additionally, a significant number of professionals display preference reversal depending on whether a decision is framed as a choice or as WTP. The magnitude of preference reversal is similar to the student sample, suggesting that information security professionals are just as susceptible to framing effects as the general population. This should concern any decision maker who would like to believe that the security recommendations they receive do not depend on the way in which they asked the question. Which type of framing tends to lead to better decisions (i.e. closer to expected value maximization) was outside the scope of our study, so more research is needed on framing effects in information security.

One behavioural anomaly which we did not find evidence for in the experiment is that information security professionals were prone to worst-case thinking. When presented with lotteries with different worst-case scenarios, professionals consistently minimized expected losses. Neither do we find evidence that decisions in our lotteries are affected when subjects are told their choices would be further evaluated. However, the lack of influence of other-evaluation on decisions may be due to a weak treatment manipulation.

Taking this evidence as a whole, we would not characterize security professionals as fully rational decision-makers. This implies that calculations involved in risk assessment methodologies are dependent on the security decision maker's subjective perceptions. This is potentially a weak link in the security chain that needs to be strengthened.

Finally, we examined security professionals' preferences between operability and security. Preferences across individuals were heterogeneous and we also find that preferences between security and operability are correlated with professional role. Professionals with more experience, and in more senior roles, tended to choose security over operability. Senior security positions tend to be associated with risk ownership and liability. Senior position also require a greater examination of the 'big picture' of the security environment. The fact that these individuals chose security over operability might indicate that professionals of such positions are more inclined to consider potentially catastrophic and disastrous outcomes that can disrupt business functions, and therefore choose the 'safer path' of security prioritisation.

References

[1] IBM Corp. Released 2012. IBM SPSS statistics for Windows, Version 21.0. Armonk, NY: IBM Corp.

[2] Alessandro Acquisti. Privacy in electronic commerce and the economics of immediate gratification. In *Proceedings of the 5th ACM conference on Electronic commerce*, pages 21–29. ACM, 2004.

[3] Alessandro Acquisti and Jens Grossklags. Privacy and rationality in individual decision making. *IEEE Security & Privacy*, 2:24–30, 2005.

[4] Ariely, Dan; Kamenica, Emir and Prelec, Dražen, 'Man's search for meaning: The case of Legos.' *Journal of Economic Behavior & Organization*, 67, 3 (2008): 671-677.

[5] Bruce Schneier. The psychology of security. In *Progress in Cryptology–AFRICACRYPT 2008*, pages 50–79. Springer, 2008.

[6] Ross Anderson. Why Information Security is Hard - An Economic Perspective. In *Proceedings of 17th Annual Computer Security Applications Conference (ACSAC)*. New Orleans, Louisiana, Dec. 10–14, 2001.

[7] Ross Anderson. Information Security Economics-and Beyond. In *Deontic Logic in Computer Science*, pages 49–49. Springer, 2008.

[8] Ross Anderson, Chris Barton, Rainer Böhme, Richard Clayton, Michel J.G. Van Eeten, Michael Levi, Tyler Moore, and Stefan Savage. Measuring the cost of cybercrime. In *The Economics of Information Security and Privacy*, pages 265–300. Springer, 2013.

[9] Ross Anderson and Tyler Moore. The Economics of Information Security. *Science*, 314(5799):610–613, 2006.

[10] Ross Anderson and Tyler Moore. Information security: where computer science, economics and psychology meet. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 367(1898):2717–2727, 2009.[11] Kenneth Joseph Arrow. *The Economics of Information (Collected Pa- pers of Kenneth J. Arrow)*, volume 4. Cambridge, Massachusetts: Belk- nap Press, 1984.

[12] Information Systems Audit and Control Association (ISACA). G41 Return on Security Investment (ROSI), 2010. Available online at www. isaca.org.

[13] Michelle Baddeley. Information security: Lessons from Behavioural Economics. Working Paper, Gonville and Caius College, University of Cambridge, 2011.

[14] S.B. 1386, 2002 Leg., Reg. Sess. (Cal. 2002).

[15] Yolanta Beresnevichiene, David Pym, and Simon Shiu. Decision support for systems security investment. In *Network Operations and Management Symposium Workshops (NOMS Wksps),2010 IEEE/IFIP*, pages 118–125. IEEE, 2010.

[16] Pedro Bordalo, Nicola Gennaioli, and Andrei Shleifer. Salience theory of choice under risk. *The Quarterly Journal of Economics*, 127(3):1243–1285, 2012.

[17] Denis Bouyssou, Didier Dubois, Henri Prade, and Marc Pirlot. *Decision Making Process: Concepts and Methods*. John Wiley & Sons, 2013.

[18] Joel Brenner. ISO 27001: Risk Management and Compliance. *Risk Management*, 54(1):24, 2007.

[19] Colin F. Camerer, George Loewenstein, and Matthew Rabin. *Advances in Behavioral Economics*. Princeton University Press, Princeton, NJ, 2011.

[20] James L. Cebula and Lisa R. Young. A taxonomy of operational cyber security risks. Technical report, DTIC Document, Carnegie Mellon University Software Engineering Institute (SEI), 2010.

[21] Clare Chua Chow and Rakesh K Sarin. Known, unknown, and unknowable uncertainties. *Theory and Decision*, 52(2):127–138, 2002.

[22] Frank Wilcoxon, S K Katti, and Roberta A Wilcox. Critical values and probability levels for the Wilcoxon rank sum test and the Wilcoxon signed rank test. *Selected Tables in Mathematical Statistics*, 1:171–259, 1970.

[23] Shawn P. Curley, J. Frank Yates, and Richard A. Abrams. Psychological sources of ambiguity avoidance. *Organizational Behavior and Human Decision Processes*, 38(2):230–256, 1986.

[24] C Derrick Huang, Qing Hu, and Ravi S Behara. An economic analysis of the optimal information security investment in the case of a risk-averse firm. *International Journal of Production Economics*, 114(2):793–804, 2008.

[25] Thomas Dohmen, Armin Falk, David Huffman, Uwe Sunde, Jürgen Schupp, and Gert G Wagner. Individual risk attitudes: Measurement, determinants, and behavioral consequences. *Journal of the European Economic Association*, 9(3):522–550, 2011.

[26] Daniel Ellsberg. Risk, ambiguity, and the savage axioms. *The Quarterly Journal of Economics*, 75(4):643–669, 1961.

[27] John Von Neumann and Oskar Morgenstern. *Theory of Games and Economic Behavior (60th Anniversary Commemorative Edition)*. Princeton University Press, 2007.

[28] ENISA. Introduction to Return on Security Investment. Technical report, ENISA, Heraklion, Greece, Dec 2012. Available online at https://www.enisa.europa.eu/activities/cert/ other-work/introduction-to-return-on-security-investment.

[29] Vilhelm Verendel. A prospect theory approach to security. Technical report, Department of Computer Science and Engineering, Chalmers University of Technology, 2008.

[30] Department for Business, Innovation and Skills (BIS, UK) and Technology Strategy Board. Cost of business cyber security breaches almost double. Technical report, April 2014. https://www.gov.uk/government/news/ cost-of-business-cyber-security-breaches-almost-double.

[31] Milton Friedman. The use of ranks to avoid the assumption of normality implicit in the analysis of variance. *Journal of the American Statistical Association*, 32(200):675–701, 1937.

[32] Vaibhav Garg and Jean Camp. Heuristics and biases: implications for security design. *Technology and Society Magazine, IEEE*, 32(1):73–79, 2013.

[33] Daniel Geer. Power. law. *Security & Privacy, IEEE*, 10(1):94–95, 2012. [34] Nicola Gennaioli and Andrei Shleifer. What comes to mind.*Quarterly Journal of Economics*, 125(4):1399–1434, 2010.

[35] Lawrence A Gordon and Martin P Loeb. The economics of information security investment. *ACM Transactions on Information and System Security (TISSEC)*, (4):438–457, 2002.

[36] Lawrence A Gordon and Martin P Loeb. *Managing Cybersecurity resources: a cost-benefit analysis*, volume 1. McGraw-Hill New York, 2006.

[37] Amos Tversky and Daniel Kahneman. Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and Uncertainty*, 5(4):297–323, 1992.

[38] Charles A Holt and Susan K Laury. Risk aversion and incentive effects. *American Economic Review*, 92(5):1644–1655, 2002.

[39] Kevin J. Soo Hoo. *How much is enough? A risk management approach to computer security*. Working Paper, Stanford University, 2000.

[40] Christos Ioannidis, David Pym, and Julian Williams. Fixed costs, investment rigidities, and risk aversion in information security: A utility-theoretic approach. In *B. Schneier (Ed.), Economics of Security and Privacy III*, pages 171–191. Springer, 2012. Proceedings of the 2011 Workshop on the Economics of Information Security.

[41] BS ISO. IEC 27005:2008. *Information Technology–Security Techniques– Information Security Risk Management*, 2012.

[42] Daniel Kahneman. *Thinking, fast and slow*. Macmillan, 2011.

[43] Daniel Kahneman and Amos Tversky. Prospect theory: An analysis of decision under risk. *Econometrica: Journal of the Econometric Society*, 47(2):263–291, 1979.

[44] Daniel Kahneman and Amos Tversky. Choices, values, and frames. *American Psychologist*, 39(4):341, 1984.

[45] Frank H Knight. *Risk, uncertainty and profit*. Courier Dover Publications, 2012.

[46] Ponemon Institute LLC. Cost of Data Breach Study: Australia. 2011. [47] Christian Locher. Methodologies for evaluating information security investments - What Basel II can change in the financial industry. 2005. In Proceedings of the 13th European conference of information systems, information systems in a rapidly changing economy, ECIS 2005, Regensburg, Germany, 26-28 May 2005.

[48] Mark J. Machina. Choice under uncertainty: Problems solved and unsolved. *The Journal of Economic Perspectives*, 1(1):121–154, 1987.

[49] Thomas Maillart and Didier Sornette. Heavy-tailed distribution of cyber-risks. *The European Physical Journal B - Condensed Matter and Complex Systems*, 75(3):357–364, 2010.

[50] Sandra Maximiano. Measuring reciprocity: Do survey and experimental data correlate. Working paper, Krannert School of Management, Purdue University, 2012.

[51] Mike McGuire and Samantha Dowling. Cyber crime: A review of the evidence. Summary of key findings and implications. Home Office Research report 75, 2013. www.gov.uk/government/uploads/system/uploads/ attachment_data/file/246749/horr75-summary.pdf.

[52] Evan Moore and Catherine Eckel. Measuring ambiguity aversion. Unpublished manuscript. Department of Economics, Virginia Tech. 2003.

[53] Mark EJ Newman. Power laws, Pareto distributions and Zipf's law. *Contemporary physics*, 46(5):323–351, 2005.

[54] Neil J. Schroeder. Using prospect theory to investigate decision-making bias within an information security context. Technical report, Dept. of the Air Force Air University, Air Force Institute of Technology, 2005.

[55] Bruce Schneier. Worst-case thinking makes us nuts, not safe. Schneier on Security (blog), May 2010. https://www.schneier.com/essay-316. html.

[56] Robert Richardson. CSI Computer Crime and Security Survey, 2008.

[57] Robert Richardson. CSI Computer Crime and Security Survey, 2010.

[58] Ron Ross, Stu Katzke, Arnold Johnson, Marianne Swanson, Gary Stoneburner, George Rogers, and Annabelle Lee. Recommended security controls for federal information systems. *NIST Special Publication*, 800:53, 2005.

[59] Michael Rothschild and Joseph E. Stiglitz. Increasing risk: I. A definition. *Journal of Economic theory*, 2(3):225–243, 1970.

A Appendix

A.1 Survey Questions

• ‘Are you related with the profession or practice of Information Security in any way?’
Yes / No

• ‘How many years of experience do you have in Information Security related tasks?’

• ‘How willing are you to take risks in general?’ *0 to 10*
0: Not willing at all 10: Very willing

• ‘Your job title most closely resembles:’
– Senior executive role (e.g. CEO, CIO, CISO, CSO etc.)
– Managerial role (e.g. Project Manager, IT Director, Security Manager etc.)
– IT & Security (e.g. Security Officer, System Administrator, Cyber Security Information Analyst etc.)
– Compliance, Risk or Privacy role (e.g. Governance, Risk and Compliance Consultant, Information Security Consultant, Auditor etc.)
– Other: *please specify*

• ‘Does your job position allow you to make independent Information Security related decisions?’ *Yes / No*

• ‘How worried are you that a severe/important security incident might materialise in your company/organisation, despite the existing protective measures?’ *0 to 10*
0: Not worried at all 10: Very worried

• ‘How worried are you about new unidentified information security threats?’ *0 to 10*
0: Not worried at all 10: Very worried

• ‘Have you experienced any important security incident in the past?’ *Yes / No*

• ‘How closely related do you think investment in Information Security is to business objectives?’ *0 to 10*
0: Not related at all 10: Very much related

• ‘How much do you think companies/organisations focus on business operations and as a result underestimate or neglect security?’ *0 to 10*

0: Not worried at all 10: Very worried

• ‘Where / to whom does your Chief Information Security Officer (CISO or CSO) or equivalent senior executive report?’

• ‘What is the size of your company?’

• ‘What is your gender?’

• ‘What is your age?’

• ‘What is your educational level?’

• ‘What is your marital status?’

• ‘What is the number of dependents in your family?’

• ‘What is your approximate annual income in British pounds?’

• ‘Which country do you live in?’

• ‘What is your nationality?’

• ‘What is your mother tongue?’

A.2 H1 Instrument

There are four types of experiment questions on willingness-to-pay to avoid a lottery, one for each lottery type. The actual values of p_i and x_i are shown in the second and third column of Table 9:

'What is the maximum amount that you are willing to pay in order to avoid playing a lottery in which there is..:

- ..a p% probability of losing \$50 and losing nothing otherwise?'
- ..a probability between p_1% and p_2% of losing \$50?'
- ..a p% probability of losing an amount between \$$x_1$ and \$$x_2$ and losing nothing otherwise?'
- ..a probability between p_1% and p_2 of losing an amount between \$$x_1$ and \$$x_2$ and losing nothing otherwise?'

Table 9: H1 Instrument

#	**Prob.** (p%)	**Outcomes** (x in \$)	**WTP**	**EV** μ	**Expected Outcome Interval**	**Outcome Range**
$\boldsymbol{H_1}$ **1**	5	-50	0 to 100	-2.5	-2.5	0
$\boldsymbol{H_1}$ **2**	0-10	-50	0 to 100	-2.5	[-5, 0]	5
$\boldsymbol{H_1}$ **3**	5	-80 to -20	0 to 100	-2.5	[-4, -1]	3
$\boldsymbol{H_1}$ **4**	0-10	-80 to -20	0 to 100	-2.5	[-8, 0]	8
$\boldsymbol{H_1}$ **5**	15	-50	0 to 100	-7.5	-7.5	0
$\boldsymbol{H_1}$ **6**	0-30	-50	0 to 100	-7.5	[-7.5, 0]	7.5
$\boldsymbol{H_1}$ **7**	15	-80 to -20	0 to 100	-7.5	[-12, -3]	9
$\boldsymbol{H_1}$ **8**	0-30	-80 to -20	0 to 100	-7.5	[-24, 0]	18
$\boldsymbol{H_1}$ **9**	50	-50	0 to 100	-25	-25	0
$\boldsymbol{H_1}$ **10**	35-65	-50	0 to 100	-25	[-32.5,-17.5]	15
$\boldsymbol{H_1}$ **11**	50	-80 to -20	0 to 100	-25	[-40, -10]	30
$\boldsymbol{H_1}$ **12**	35-65	-80 to -20	0 to 100	-25	[-52, -7]	45

A.3 Lottery Comparisons for Hypothesis 2

H_2 1 (Hypothesis 2 Question 1): comparison of Lotteries 5 and 6

H_2 2 (Hypothesis 2 Question 2): comparison of Lotteries 6 and 7

H_2 3 (Hypothesis 2 Question 3): comparison of Lotteries 4 and 2

H_2 4 (Hypothesis 2 Question 4): comparison of Lotteries 2 and 3

H_2 5 (Hypothesis 2 Question 5): comparison of Lotteries 1 and 8

Lottery 1
a probability of 85% of 50 a probability of 8% of losing 150 a probability of 3.5% of losing 300 a probability of 2.5% of losing 450 a probability of 1% of losing 1000
μ = -86.25, Var = 14698.4

Lottery 2
a probability of 15% of losing nothing a probability of 30% of losing 166.66 a probability of 30% of losing 300 a probability of 20% of losing 450 a probability of 5% of losing 900
μ = -274.998, Var = 40708.8

Lottery 3
a probability of 15% of losing nothing a probability of 30% of losing 183.33 a probability of 30% of losing 300 a probability of 20% of losing 450 a probability of 5% of losing 800
μ = -274.999, Var = 33958.5

Lottery 4
a probability of 15% of losing nothing a probability of 30% of losing 200 a probability of 30% of losing 300 a probability of 20% of losing 450 a probability of 5% of losing 700
μ = -275, Var = 28375

Lottery 5
a probability of 85% of losing 45 a probability of 8% of losing 220 a probability of 3.5% of losing 300 a probability of 2.5% of losing 450 a probability of 1% of losing 900
μ = -86.6, Var = 14406.2

Lottery 6
a probability of 85% of losing 50 a probability of 8% of losing 170 a probability of 3.5% of losing 300 a probability of 2.5% of losing 400 a probability of 1% of losing 1000
μ = -86.6, Var = 14087.4

Lottery 7
a probability of 85% of losing 45 a probability of 8% of losing 250 a probability of 3.5% of losing 350 a probability of 2.5% of losing 450 a probability of 1% of losing 800
μ = -89.75, Var = 14416.2

Lottery 8
a probability of 85% of 46 a probability of 8% of losing 180 a probability of 3.5% of losing 350 a probability of 2.5% of losing 480 a probability of 1% of losing 900
μ = -86.75, Var = 15012.5

A.4 Willingness-to-pay Lotteries for Hypothesis 2

H_2 6 (Hypothesis 2 Question 6): WTP for Lottery 5

H_2 7 (Hypothesis 2 Question 7): WTP for Lottery 6

H_2 8 (Hypothesis 2 Question 8): WTP for Lottery 7

Definitions

$H_x y$	A lottery with index *y*, that is mainly related to hypothesis *x*.
$H_1 1$ *to* $H_1 12$	Two-outcome lotteries with negative or zero outcomes; participants stated their willingness-to-pay to avoid these lotteries.
L_i	Various five-outcome lotteries used in lottery comparisons.
Group A	Lotteries $H_1 1$ to $H_1 4$ with expected value $\mu = -2.5$.
Group B	Lotteries $H_1 5$ to $H_1 8$ with expected value $\mu = -7.5$.
Group C	Lotteries $H_1 9$ to $H_1 12$ with expected value $\mu = -25$.
Scenario1	Experiment question in which participants chose between enhancement of either security or operability.
Scenario2	Experiment mechanism in which participants chose between: A) remaining in the current system state, B) enhancement and reduction of security and operability (based on previous answers) and C) indifference between A and B.