Experimental Elicitation of Risk Behaviour amongst Information Security Professionals

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

[1]Information Security Group, Royal Holloway, University of London, UK
Konstantinos.Mersinas.2011@live.rhul.ac.uk ,
Keith.Martin@rhul.ac.uk
[2]Department of Economics, Royal Holloway, University of London, UK
{Bjoern.Hartig, A.Seltzer}@rhul.ac.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 unbiased and rationally, we conducted an economic online experiment and survey measuring risk attitude of security professionals and contrasting their behaviour with the general population. Participants were asked to state their willingness-to-pay in order to avoid a series of losses-only lotteries and to make choices between such lotteries. We also devised a mechanism to elicit preferences between security and operability. Our findings suggest that security professionals are risk and ambiguity averse, consider small losses inevitable and take risks when losses are associated with large probabilities. We find that their preferences are measurably different from those of the general population in some of these aspects. We also find that job position influences security and operability preferences and that avoidance of salient (catastrophic) outcomes explains some of the professionals' behaviour. Moreover, professionals are susceptible to framing effects to the same extent as the general population, and reveal distorted probability perception, factors that are usually overlooked in risk assessment methodologies.

1 Introduction

Spending on protective measures and mechanisms for information security is a big issue for most organizations. Specifying the optimal level of information security investment is not an easy task for security professionals. Reports

show that more and more money is invested on defensive security measures and there are indications that the cost of security breaches either remains at high levels [51] or has a growing tendency [30, 46]. Despite the fact that the budget for security investment is increasing, insufficient expenditure on information security is considered as one of the main obstacles that security professionals face [30], making security investment level optimisation crucial, but hard to achieve, balancing between overspending and insecurity [1].

Cost-benefit analysis [35] as well as risk-management approaches [39] constitute a widely accepted solution to tackle this problem and there is a variety of models used by professionals: Net Present Value (NPV), Internal Rate of Return (IRR), Return of Investment (ROI) and Return of Security Investment (ROSI) [12, 28, 36, 47, 57].

It is however clear that the field of economics has not provided a dominant model for decision-making in security [56] and professionals are encouraged to choose their own appropriate risk analysis and assessment methods [18,41] to match the needs of their organisations.

However, all quantitative risk assessment methodologies are subject to three significant limitations [28]:

1. they are based on many approximations (e.g. unknown risks);
2. these approximations are often biased by the perception of risk, and;
3. the involved calculations can be easily manipulated.

Subjectivity of risk perception and the lack of a predominant bounding economic model for *deciding* and *justifying* security investment, signify the importance of the decision-maker's preferences and risk attitude. The field of behavioural economics [19,44,48] reveals various heuristics and biases that individuals use when making decisions.

For example, a security professional has to decide the amount of protective investment that has to be spent in order to *avoid* unwanted losses. She might posses data on past incident occurrences, but it is up to her to decide and propose the exact investment level. In such a scenario, her attitude towards risk can be differentiated depending on the probability of a threat materialising, and also on the expected damage to be incurred. When a threat bears potential catastrophic outcomes, the attention of the professional might be disproportionately focused on the most *salient*, worst-case

[1] Cybersecurity survey data should be carefully interpreted, as contacted and responding populations can lead to unrepresentative samples [29]. Also, surveying rare events is by default problematic.

outcome, and hence she might be willing to spend more in order to be on the safe side, even if the probability of such an event is negligible. In other cases, 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, another factor that potentially affects her decision. Finally, the level of investment has to be communicated to other parties in the organisational structure; these parties may lack the expertise necessary to understand her suggestions. Such a possibility might cause the security professional to either exaggerate or deflate her initial proposals for making them justifiable.

Researchers have used behavioural theories, such as prospect theory [43, 64], in the context of information security [62, 65] but, to our knowledge, actual behaviour between and within security professionals has been relatively less studied. Our study contributes in understanding the attitude of active information security professionals and practitioners across various levels of risk and uncertainty and in comparing the behaviour of professionals against the behaviour of the general population [2]. We expect that experienced professionals are better at estimating expected outcomes of lotteries than the sample of students. A clear understanding of potential behavioural biases can constitute a handy tool for decision-makers as it can lead to the development of appropriate strategies for mitigating (or amplifying) the relevant biases.

For the purposes of eliciting risk attitudes from security professionals we abstract potential vulnerabilities (probabilities) and losses (outcomes) in the form of lotteries. We also use scenarios with information security context that involve defence costs and direct losses, in the spirit of [8], as well as operational losses.

Moreover, the environment of information security has inherent characteristics that diversify the context of decisions. In particular, we focus on the following distinctive set of features, which are hypothesised and examined in our experimental approach:

1. *Loss domain*: each security investment decision can be described as a lottery with losses only. The best outcome is zero, so that the scope of the decision-maker is *loss prevention.*

[2]We consider a sample of students randomly drawn from the database records of the Laboratory for Decision Making & Economic Research at Royal Holloway University of London, in order to contrast with behaviour of professionals. These are students that come from all departments and faculties of the university. We use the terms 'general population' and 'student sample' interchangeably.

2. *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.

3. *Security and operability*: in each decision there is an inherent trade-off between security and operability of the system, with both having measurable monetary costs.

We find that *security* versus *operability* preferences of professionals are significantly diversified across different job positions and that professionals' risk attitude is distinguishable from the behaviour of the student sample. Also, although professionals can make more accurate predictions than the general population, they have, to a certain degree, a distorted understanding of probabilities.

The rest of the paper is organised in the following way. In Section 2 the theoretical framework of behavioural economics is presented, along with the economic theory behind the experimental and survey empirical approach. In Section 3 the core hypotheses and the experiment design are presented in detail. Findings of data analysis are listed in Section 4. We discuss findings in Section 5 and conclude in Section 6.

2 Background and Approach

We use an economic experiment for risk attitude elicitation, and a self-reported survey, in order to examine potential behavioural biases of security professionals and to contrast their risk attitude against behaviour of the general population. The theoretical framework for our empirical approach is the field of behavioural economics, and in particular behavioural approaches in information security.

2.1 Behavioural Economics and Security

The study is motivated by the general question of 'how much security is needed', i.e. the estimation of the appropriate level of security investment. However, our approach focuses on the human factor, the risk perception of the decision-makers themselves. The importance of the economics of information security with extensions to behavioural aspects has been highlighted in various papers of Anderson and Moore [6, 7, 9, 10]. More broadly, behavioural economics have revealed a number of 'paradoxes' or systematic violations of expected utility theory [66] showing that the rational-agent 'homo economicus' is not observed empirically [19, 44, 48]. Bruce Schneier sketches the effect of heuristics and biases as the psychology of security, describing risk and uncertainty perception issues [60].

There have been studies that use an expected utility theory [24] as well as prospect theory approach in security [65]. Schroeder uses prospect theory and also introduces the dichotomy between security and operations in a military-context empirical research [62]. Insights from psychology and sociological factors as well as biases in security are presented by Baddeley [13] and there has been focus on the decision-making process of security professionals from a decision support system point of view [15]. Shiu et al. conducted an experiment on security professionals with economic framing controls, revealing the existence of the confirmation bias [14]. Other biases, like the status quo and present bias have been specifically targeted, even if from a privacy perspective [2–4]; the effect of biases on security design has also been explored [32]. Timing preferences about security investment have been studied by Ioannidis et al. [40].

Kahneman and Tversky proposed the four-fold pattern of risk attitudes [43], which states that decision-makers are risk-averse for small-probability losses and large-probability gains and risk-seeking for small-probability gains and large-probability losses. The pattern has been adopted and studied in various environments [37], but it is not known whether it fits analogously in the information security losses-only context.

Salience theory [16] states that it is salience of outcomes, instead of the probabilities, that attract the focus of the decision-maker. Salience is the phenomenon in which "when one's attention is differentially directed to one portion of the environment rather than to others, the information contained in that portion will receive disproportionate weighting in subsequent judgements" ([63]). Worst-case is a frequently used term in this study and it can be considered as an unusual case: "our mind has a useful capability to focus on whatever is odd, different or unusual" ([42]). Salience is formalised by ordering and diminishing sensitivity, and is therefore in accordance with rank-dependent models of choice, prospect theory in particular. We use salience theory in the experiment design to examine whether ranking of lottery payoffs, as well as payoff-magnitude, influence decisions. A lottery's worst payoff is also salient in case that its absolute difference from the rest of the outcomes, is 'sufficiently' big.

Various psychological sources have been proposed regarding ambiguity aversion [23]. According to the other-evaluation hypothesis, decision-makers become ambiguity averse when they anticipate evaluation of their choices by others, e.g. by peers, colleagues and so on. An explanation is that the individual chooses the most *a posteriori* justifiable option. A decision is more sound if it is less risky, and 'riskiness' can be measured by the disfavour of mean preserving spreads of a lottery [59]. Such behaviour can be measured with the risk premium (or certainty equivalent) that the decision-maker is

willing to pay for not taking the lottery. Other-evaluation exists in organisational contexts, as security professionals propose and justify solutions to business managers or higher management. Professionals are predicted to seek an ex post justification of their choices which places them in a defensive position to prove soundness of their decisions.

2.2 Experimental elicitation of Risk Attitudes

After three decades of economics research on decision-making under risk and uncertainty of individual choice, and especially by connecting empirical results with theory, there have been quite a few behavioural models developed. The von Neuman-Morgenstern expected utility theory axiomatisation (1947, [66]) approached decision-making from a normative (idealised decision-maker) and prescriptive (practical directions) model view. A variety of studies have indicated violations of rationality axioms in practice.

A common formalisation used in behavioural models is the notion of a *lottery* (or *prospect*), a "list of consequences with associated probabilities". It is practical to consider a finite number of possible states that completely determine a finite number of lottery consequences. If the consequences are ordered then they can be considered as utilities that the decision-maker has to choose from or maximise. This is the 'rational model of choice under uncertainty' (Arrow, [11], p60) and it can be assumed that the individual can assign subjective probabilities in each state of nature.

In behavioural models there is a distinction, originally established by Knight [45], between *risk*, in which outcomes and probabilities of lotteries are known, and *uncertainty*, where at least some of the outcomes or probabilities are unknown. Risk is 'a quantity susceptible of measurement', whereas uncertainty is 'immeasurable risk' ([45]) and *ambiguity* is the 'uncertainty about probabilities' ([26]). The concept of *ambiguity avoidance* or *ambiguity aversion* can be loosely defined as the attitude of preferring risky lotteries over uncertain lotteries, i.e. preferring known probabilities over unknown ones.

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). It should be also noted that all lotteries in the experiment are *decision-based*, with no feedback given after a choice is selected.

Willingness-to-pay (WTP) is the maximum amount that the individual is willing to sacrifice in order to avoid an undesirable event. We use WTP as a technique to model choices in the experiment, asking professionals to avoid

lotteries with only negative outcomes. We create a new instrument for measuring risk and ambiguity aversion, as a modification of the Holt and Laury instrument [38] and similar to the alternative of Moore and Eckel [52]. Some studies use outcome-ambiguous lotteries [27], while others use probability-ambiguous lotteries [5]. Our approach uses sets of lotteries with different levels of expected losses, in each of which there are four lotteries spanning from risky lotteries to lotteries ambiguous in probabilities, in outcomes and in both probabilities and outcomes. This design allows for between-subjects, as well as within-subjects analysis across lotteries of the same expected value.

Laboratory experiments are susceptible to low incentives and therefore to unrealistic results regarding risk measurement. We overcome this issue by presenting simple choice tasks, in a similar fashion to [38]. Additionally, we provide performance-dependent payment to the participants, based on their lottery replies.

2.3 Surveys

Expanding on the aforementioned core hypotheses, this study enhances the accuracy of experimental results by combining them with survey data. In general, data produced by experiments in a controlled fashion is considered more reliable, mainly because it is elicited from incentivised participants. Survey data, on the other hand, might amplify the effects of misunderstanding of questions, of information recalling or might be influenced by socially acceptable answers. Note that experiments may be free from such measurement errors, but can also be immune to certain biases, e.g. as is observed when people respond differently to hypothetical than to real situations or when they reply as if they were another person [50].

One question that follows naturally is whether observed behaviour during the experiment is correlated with the self-reported answers to the survey questions. A study that gives strong evidence supporting the validity of survey results is [25] in which risk attitudes were accurately depicted both by survey data and experiment input.

Indicatively, the survey involves questions on bad experiences from serious information security incidents, about how worried the subjects feel regarding existing and unknown security threats, and so on. We examine correlations of responses to such questions with observed risk attitude throughout the experiment. All survey questions are included in Appendix A.4.

3 Methodology

3.1 Research Questions

Through an online experiment and a complimentary survey, we analysed behaviour of security professionals on the following hypotheses:

1. *Security professionals exhibit risk and ambiguity aversion*: for a given lottery with small probabilities of losses, we predicted that professionals are willing to pay more than the expected value, in order to avoid playing that lottery. This expectation should be amplified when instead of a specified probability there are a range of probabilities, or a range of outcomes, or both (different types of ambiguity). The phenomenon is expected to be inverted for larger probabilities revealing risk-seeking behaviour, assuming that the four-fold pattern of risk attitude holds [43].

2. *Security professionals exhibit worst-case thinking*: it is expected that lotteries with salient or potentially catastrophic outcomes attract the attention of the decision-makers more than lotteries with moderate outcomes [16], and increase willingness-to-pay (WTP), irrespectively of the expected value or the variance of the lottery.

3. *Other-evaluation hypothesis*: under this hypothesis [23], it is predicted that when decisions are revealed to other parties for evaluation, individuals tend to be more ambiguity averse.

4. *Security and Operability*: we expected security professionals to show a tendency to highlight the importance of security and underweight the need for operability.

In terms of targeted populations, we conducted the same experiment with two different samples. Current and previous students of the distance learning MSc in Information Security offered by Royal Holloway, University of London (RHUL) were contacted as well as individuals registered in the experiments database of the Laboratory for Decision Making and Economic Research at RHUL. The former sample (referred to as 'professionals') mostly contains security professionals that work in the industry and decided to undertake the distance learning master's program on a part-time basis. The mean of the professionals' industry experience is $\mu = 8.95$ years. The latter sample (referred to as 'students' or 'general population') mostly consists of active full-time students that come from all departments and faculties of the university. The survey question that diversified security professionals from the 'general population' was whether subjects are 'related to the profession or practice of information security' (see Appendix A.4).

Lotteries used in the experiment were abstracted so that context-free risk attitude was observed. The risk attitude elicitation instrument was designed to capture different types of ambiguity across three levels of expected losses and three levels of probabilities. This way between-subjects WTP was measured for both populations, but also within-subjects attitude across different types of risky and ambiguous decisions is observed. In particular, a question of interest was how risk perception changes when both vulnerability (probability) and loss (outcome) are changed one at a time, or simultaneously [3].

A challenging point of the design was the creation of five-outcome lotteries for testing the worst-case thinking hypothesis. The variables that are changed across the lotteries are best-outcome, worst-outcome, expected value and variance. Moreover, certain lotteries were built on power-law distributions, as it has been shown that occurrences of many natural and social catastrophic phenomena follow such distributions [53]. Worst possible outcomes are deemed salient only if they are significantly different from the rest of the choice context, otherwise their associated events can be underweighted instead of overweighted by the participants. This means that both ranking and magnitude of losses are important. The degree of distortion of the perceived probabilities was estimated by salience theory assumptions.

Information security managers and decision-makers have to justify their investment proposals to business managers, chief officers, the board of directors or a similar body. The other-evaluation hypothesis as defined by Curley et al. [23] states that: "a decision maker, in making a choice, anticipates that others will evaluate his or her decision; and, so, makes the choice that is perceived to be most justifiable to others. This choice is for the option having the smallest degree of ambiguity". The hypotheses aimed to reveal evaluation by others as a possible psychological source of behaviour that directly influences investment choices.

Testing the other-evaluation hypothesis was ambitious in the context of an online experiment, because we had to find a way to provide an impression of an additional evaluation, on top of the standard statistical analysis that subjects know they were being subjected to.

Finally, security professionals and practitioners supposedly have a tendency to pay attention to security issues at the expense of operational issues

[3]Participants were informed that they would receive a fixed participation payment and an additional potentially larger amount depending on their 'performance'. 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.

that could be important from a business perspective. Security and operability are presented to subjects in a realistic scenario. To make the distinction clear, from potential operational risks [20] operability was framed as the *operational time* needed for task completion, and was measured explicitly in monetary terms, as was security. To exclude other factors, the scenario described an information system of moderate-impact to confidentiality, integrity and availability [58]. The experiment design measured not only the actual preference between security and operational time, but also the *relative loss aversion* in security and operability, by a series of questions dynamically linked to the subjects' previous replies.

3.2 Design

3.2.1 Hypothesis 1: Security professionals exhibit risk and ambiguity aversion

The instrument for the first hypothesis consists of 12 lotteries. Subjects were asked to state their WTP (or certainty equivalent) in order to avoid each lottery. All outcomes were in the domain of losses. The actual lottery values are depicted in the second and third column of Table 9 (Appendix A) and, for example, for lottery 1 the question presented was 'What is the maximum amount that you are willing to pay in order to avoid playing a lottery in which there is a 5% probability of losing $50 and losing nothing otherwise?'. For the purposes of risk attitude elicitation, all lotteries were context-free, i.e. abstracted from any relation to information security.

3.2.2 Hypothesis 2: Worst-case thinking

This part of the experiment consists of five pairwise lottery comparisons (Appendix A.2) for which subjects were asked to choose their preferred lottery. All lotteries consisted of five outcomes; for conforming with salience theory [16], probabilities were kept the same in both lotteries, whereas outcomes are different, so that the expected value was the same in some pairs and different in others. For three of the lotteries involved in the comparisons there was a subsequent WTP question (Appendix A.3) similarly to the instrument of Hypothesis 1 (Appendix A.1). Thus, consistency of replies could be checked between comparisons and WTP per lottery.

For example, if μ_i is the expected value of lottery L_i, Var_i its variance, and '$\succeq$', '$\succ$' denote *weak* and *strict preference* respectively, then for lotteries 9, 10 and 11 (Appendix A.3), theory predicts that:
$L_{10} \succeq L_9$, as $\mu_{10} = \mu_9$ and $Var_{10} < Var_9$
$L_{10} \succ L_{11}$, as $|\mu_{10}| < |\mu_{11}|$ and $Var_{10} < Var_{11}$
$L_9 \succ L_{11}$, as $|\mu_9| < |\mu_{11}|$ and $Var_9 \approx Var_{11}$
So, the worst lottery is L_{11}, the least damaging is L_{10}, and L_9 lies in-between:

$L_{10} \succeq L_9 \succ L_{11}$ for an expected utility maximiser.

Focusing on the worst scenarios is a purported phenomenon that the security community is aware of and which has been generally described [61]. However, there has been no attempt to quantify the manifestation of the phenomenon. The approach taken here, examines the effect of salient outcomes on the decision-maker's perception [16]. Losses in security can be catastrophic and salience is used as a means to describe attention focus on worst outcomes; subsequently, the magnitude of probability distortion is estimated. Some lotteries in the experiment are designed to approximate power-law distributions. Such distributions simulate the occurrence of rare events that are observed in various physical and social phenomena, from earthquakes to citations and web hits [53]. Moreover, there is evidence for the existence of power-laws in cyber risks and the growth of networks, relating these distributions with security issues like identity theft and malware spreading [33, 49]. In a general form of a power-law distribution, probability p is specified as a function of outcome x: $p(x) = \frac{\kappa}{(-x)^{\alpha}}$, where α is the distribution exponent and κ a constant. A rough requirement that is sustained by goodness-of-fit of various empirical data to such distributions [22] is that, $\alpha \in (0, 3)$. For the purposes of our experiment, and in order for the discrete distributions of monetary losses to approximate a power-law distribution, we have set $\alpha = 1.1$, constant $\kappa = 20$ and $x \in [-1000, 0)$.

3.2.3 Hypothesis 3: Decision-makers exhibit the 'other-evaluation' ambiguity aversion

This hypothesis was tested in the experiment by creating the following treatment: subjects were randomly divided into two groups. The first group, the *control group*, was presented with the standard version of the experiment for all hypotheses. The *treatment group* on the other hand, was initially informed that all choices made in the experiment would be 'further viewed' and would 'go through an additional evaluation process'. The experiment was controlled for missing information effects: instructions were clear and communicated to participants that there was no additional, hidden information that was known to the experimenters/evaluators and which was kept secret from the participants. Such suspicion of information asymmetry has been argued to cause more conservative behaviour in experiments [21]. The statement regarding the implied evaluation process was:

"**Important note**: Your choices and their corresponding possible outcomes in the following experiment **will be further viewed** and will go through an **additional evaluation process**, after the completion of the experiment."

The statement was deliberately left vague, as its purpose was merely to create the appropriate framing for the treatment group.

3.2.4 Hypothesis 4: Decision-makers overweight the need for security and underweight the importance of operations

This experiment phase was divided in two parts. The first part elicited preferences between enhancing security and enhancing operability of the system. It consisted of scenario-based questions in which the participants had to choose between measures A and B, where A and B had different impact on the security level and the operability of the system. Both attributes had equal monetary values assigned to them:

'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 measures do you prefer?'

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

Subsequent questions were formed dynamically, depending on previous answers, so that in the next question the value of the preferred measure was marginally decreased whereas the value of the measure that was not chosen was maintained. The sequencing was repeated until the subject crosses over from choosing one measure to another, so that a switching point between security and operability was specified.

The second part of this phase provided a measure for relative loss aversion on the attribute that was chosen in Scenario 1 (security or operability). Before Scenario 2 was presented, other questions unrelated to the hypothesis under examination were inserted in the experiment flow. These questions acted as filters to diminish the relation between the two scenarios in the participants' perception. Subjects were then asked to choose between three measures (Table 2):

Table 2: Scenario 2 template question

Choice A	Mechanism 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 Mechanism B constitute the switching point that was elicited from Scenario 1. The following questions were again shaped according to the participant's choice. If the answer was B or C, i.e. the participant chose the new proposed state or was indifferent between measures A and B, then this stage of the experiment ended. If the answer was A, then the next question was identical, except for the value of security reduction x, which was reduced to $(x-1)\%$. Note that in this example the subject is assumed to have chosen security over operability in the initial question; should the subject choose operability in the initial question then x would denote operability reduction and y security enhancement. The process continued until the reduction value reached 0 or Mechanism B or Choice C was chosen at any point.

The difference i between value x of Scenario 1, and the final value $x-i$, $i=0,1,..,$ as presented in Mechanism B in Scenario 2, is the magnitude of *loss aversion* on the preferred attribute (security or operability).

We assume that utility functions of security and operability are $Sec(.)$ and $Ops(.)$ respectively. The utility functions are defined on $[-1,1] \longrightarrow \mathbb{R}$, and also $Sec(a)$ and $Ops(a) > 0$ iff $a > 0$, $Sec(a)$ and $Ops(a) < 0$ iff $a < 0$ and $Sec(a)$ and $Ops(a) = 0$ if $a = 0$.

For example, assuming that the switching point elicited from Scenario 1 was $Sec(x)$ and $Ops(10)$, $x \in [0,9]$, then we can assume for simplicity [4] that

$$Sec(+x\%) = Ops(+10\%). \tag{1}$$

Assuming that in Scenario 2, the current state A was preferred to Mechanism B:

$$Sec(-x\%) + Ops(+10\%) < Sec(0) + Ops(0) = 0 \tag{2}$$

$$\Rightarrow -Sec(-x\%) > Ops(+10\%) \tag{3}$$

$$\Rightarrow -Sec(-x\%) > Sec(+x\%). \tag{4}$$

Inequality (3) implies that the individual manifests *relative loss aversion* between the two attributes (security and operability), as $x \in [0,9]$, and Inequality (4) that there is *loss aversion* on the utility of the preferred attribute (here on security). By the assumed utility functions we see that the absolute value of the utility of a reduction is greater than the utility of an enhancement of the same value. In other words, a reduction 'hurts more' than an enhancement satisfies.

[4] It would be more precise, e.g. for $Sec(5\%) < Ops(10\%) < Sec(6\%)$, to have an approximation of $Ops(10\%) = Sec(\zeta\%)$, $\zeta \in (5,6)$.

If Mechanism B had been chosen in the initial question of Scenario 2, this would mean that

$$\begin{aligned} Sec(0\%) + Ops(0\%) &< Sec(-x\%) + Ops(10\%) \\ \Rightarrow Ops(10\%) &> -Sec(-x\%) \\ \Rightarrow -Sec(-x\%) &< Sec(x\%). \end{aligned} \tag{5}$$

Therefore, no relative loss aversion is manifested between the attributes or on the attribute of security. Quite the contrary: enhancement is preferred to reduction, so reduction 'hurts less' than enhancement.
If Mechanism B was chosen in subsequent questions of Scenario 2, then e.g.

$$\begin{aligned} Sec(0\%) + Ops(0\%) &< Sec(-(x-1)\%) + Ops(10\%) \\ \Rightarrow Ops(10\%) &> -Sec(-(x-1)\%), \end{aligned}$$

which also does not imply any loss aversion. However, if Mechanism A was again chosen in the first subsequent question of Scenario 2, then

$$\begin{aligned} Sec(0\%) + Ops(0\%) &> Sec(-(x-1)\%) + Ops(10\%) \\ \Rightarrow Ops(10\%) &< -Sec(-(x-1)\%) \\ \Rightarrow -Sec(-(x-1)\%) &> Sec(x\%), \end{aligned} \tag{6}$$

which would mean that the magnitude of loss aversion is increased in Inequality (6) in comparison to Inequality (5). This magnitude is captured in the variable $LOSS\_AV\_SEC$ for individuals that initially preferred security and similarly in $LOSS\_AV\_OPS$ for operability (Figures 16, 17). So, an observed value of loss aversion κ, say in security, is translated as $-Sec(-\kappa\%) > Sec(+\lambda\%)$ or $|Sec(-\kappa\%)| > |Sec(+\lambda\%)|$, with $\kappa, \lambda > 0$ and $\kappa \in (0, \lambda]$ [5].

If Choice C was chosen in the initial question of Scenario 2, this would mean that

$$\begin{aligned} Sec(0\%) + Ops(0\%) &= Sec(-x\%) + Ops(10\%) \\ \Rightarrow Ops(10\%) &= -Sec(-x\%) \\ \Rightarrow Sec(x\%) &= -Sec(-x\%). \end{aligned} \tag{7}$$

That is, preferences would be linear. Therefore, no further actions need to be taken in case of choices B or C.

The 'Display Logic' diagram that was used for the design of this experiment phase in the Qualtrics software [54] is presented in Figure 1.

[5] The last pair of the utilities $Sec(.)$ and $Ops(.)$ that can possibly be compared by participants is $Sec(-1\%)$ against $Ops(+10\%)$ or vice versa. So, if the last choice is still the current state, the final (and maximum) loss aversion score is 9. With the current simplification in the formalism, this means $Sec(-\epsilon\%)$, with $\epsilon > 0$ and ϵ very small; ϵ cannot be interpreted as $\epsilon = 0$, as it was initially assumed that $Sec(0) = Ops(0) = 0$.

Figure 1: Display Logic diagram for Hypothesis 4.

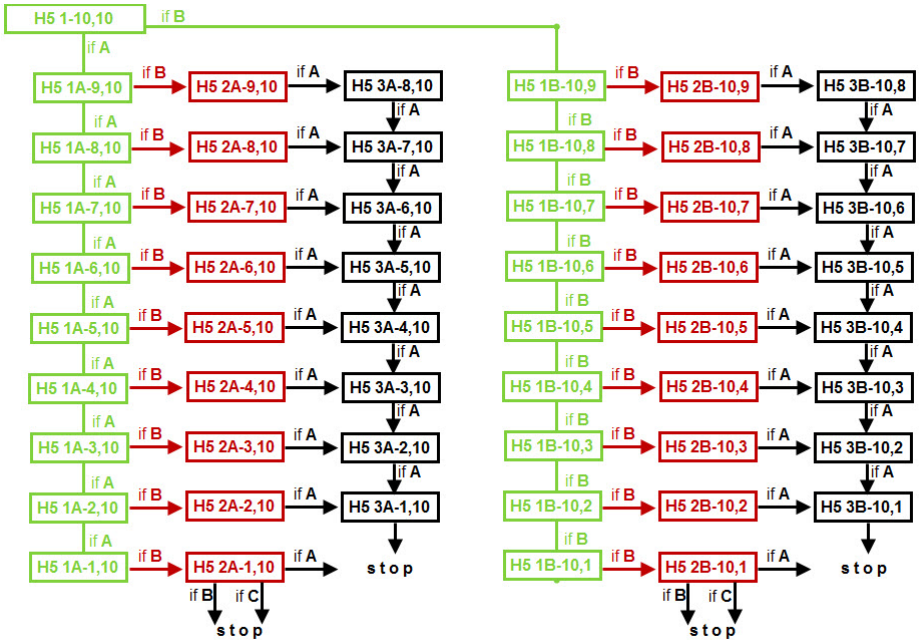

Each box represents a question in the experiment. H5 is the coding used for this series of choices; '1A' denotes a preference for security in Scenario 1 and '1B' for operability. This first series of questions is used to trace the flip point where preference changes from security to operability or vice versa, and is stored as variable 'H5 2'. In the questions with coding 'H5 2' the percentage value of the switching point is depicted (e.g. (9,10) indicates $Sec(9\%)$ and $Ops(10\%)$), and this pair is subsequently presented in a three-choice question of Scenario 2. Finally, the 'H5 3' questions serve the purpose of gradual reduction of security or operability (coding 3A and 3B respectively) of Scenario 2, whenever choice A ('Remains at the current system state') is selected. The process is terminated if choices B or C are selected at any point.

4 Analysis and Findings

The experiment was designed in such a way that there was control for potential *order effects* (Appendix B: Order Effects B.3). Data was checked for validity and *cleaned* accordingly (Appendix B: Data Cleaning B.1) and *outliers* have been shown to be non-influential (Appendix B: Outliers B.2) for the relevant tests (e.g. the one-sample t-test in Section 4.1.1).

The following sections present the findings along with their methodology and implications.

4.1 Risk and Ambiguity Aversion

The main hypothesis of risk and ambiguity aversion was examined both amongst independent subjects and per subject. The following lottery categorisation was used for both between- and within-subjects tests and its purpose was to examine whether the magnitude of losses and the nature of stakes (risky or ambiguous or both) have effects on the WTP of the participants.

- *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 $\mu = -25$.

Group A corresponds to the first four lotteries of Table 9 in Appendix A (H1 Instrument), Group B consists of lotteries 5 to 8, and the last four lotteries of the table are in Group C. It should be noted that the first lottery of each group is a risky lottery, that is, it contains specific probability and outcome values. The second lottery of each group has specific losses and a probability interval, i.e. it is probability-ambiguous. The third lottery of each group is outcome-ambiguous, and the last lottery of each group is both probability- and outcome ambiguous.

Findings on Risk Aversion

Finding 1: Both professionals and students are risk averse for small probability losses, but become risk seeking for very likely losses.

Both security professionals and students were willing to pay significantly more than the expected value of the first eight lotteries (Groups A and B) of Table 9, which included both risky and ambiguous lotteries. In particular, significant risk aversion was manifested in the lotteries with small ($p = 0.05$) and medium ($p = 0.15$) actual or average probabilities, which correspond to small ($2.5) and medium-range ($7.5) expected losses. Table 3 in Section 4.1.1 depicts mean differences between stated WTP and expected value of each lottery, for both samples. In other words, Table 3 reveals the lotteries for which WTP of subjects was significantly different from the expected loss.

However, both security professionals and students became risk seeking when the probability of loss was large, switching from their risk averse behaviour exhibited in the first eight lotteries (Figure 2). In the experiment 'large' probability was manifested as $p = 0.5$. The detailed methodology and analysis for these results are presented in Section 4.1.1.

Figure 2: Mean Risk Averse (positive) and Risk Taking (negative) WTP of Students and Professionals per lottery.
Bars represent participants' mean WTP minus the EV of each of the 12 lotteries.

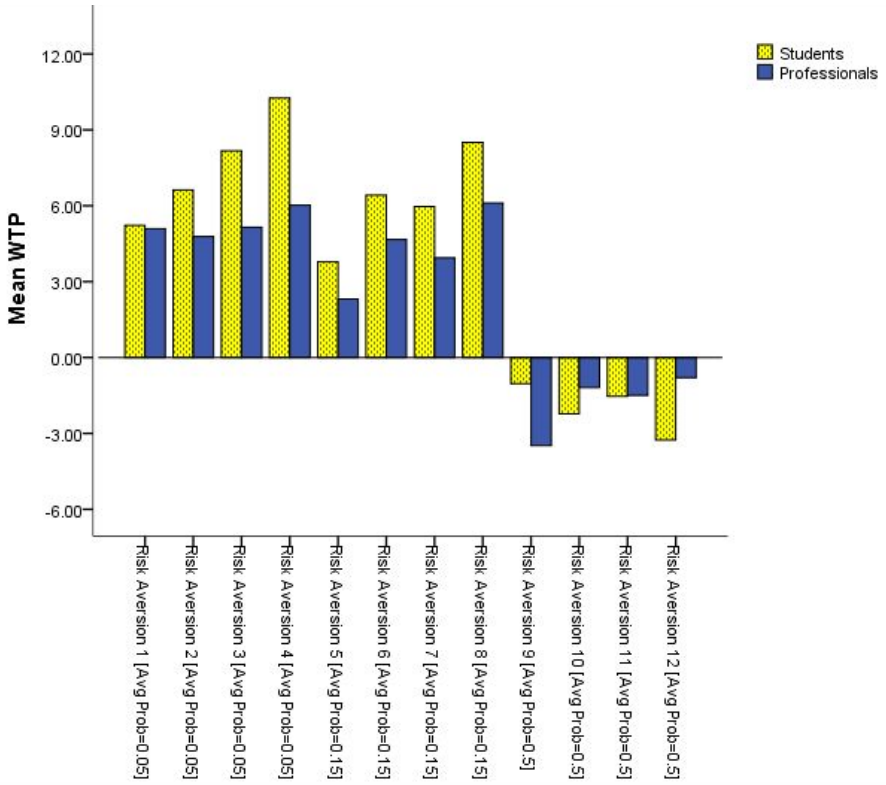

Findings on Ambiguity Aversion

Finding 2: Professionals reveal ambiguity aversion in all of their choices; such aversion is not consistently observed for the general population.
Finding 3: Professionals are better at estimating expected losses than the general population.

Security professionals became more risk averse when they confronted ambiguity, compared to when they confronted risk. This result does not hold for the general population in all cases.
We considered how WTP changes *within*-subjects, i.e. how each subject diversifies its WTP when presented with different types of risky and ambiguous lotteries. Figures 4, 5, 6, 7 and 8 show differences between risky and ambiguous lotteries within subjects of both samples. In all three groups (A, B and C) professionals revealed significant differences in WTP between at least two lotteries of each group; differences were revealed, as expected, in the pair of each risky lottery with the lottery that was ambiguous in both

probabilities and outcomes. Students did not reveal significant differences amongst the lotteries of all groups. Detailed methodology and analysis for these results are presented in Section 4.1.2.

It is notable that no statistically significant diversification was observed amongst the two types of ambiguous lotteries: the lotteries with ambiguous probabilities and the lotteries with ambiguous outcomes.

However, professionals' WTP was closer to the expected value than the general population. Specifically, professionals' WTP had smaller mean difference from the test value (zero), i.e. from the lottery expected value, than the WTP of students in 11 out of the 12 lotteries (Table 3). Remarkably, the only lottery in which professionals on average were willing to pay an amount that was more distant from the expected value than the students' amount was lottery $H_1 9$, with a large loss probability $p = 0.5$; in the consequtive lotteries risk seeking attitude was observed. This means that, in general, professionals' estimations are closer to the expected value than students' WTP.

Figure 3: Interaction of *Pro or Student* and $H_1 9$ with *General Risk* as moderator

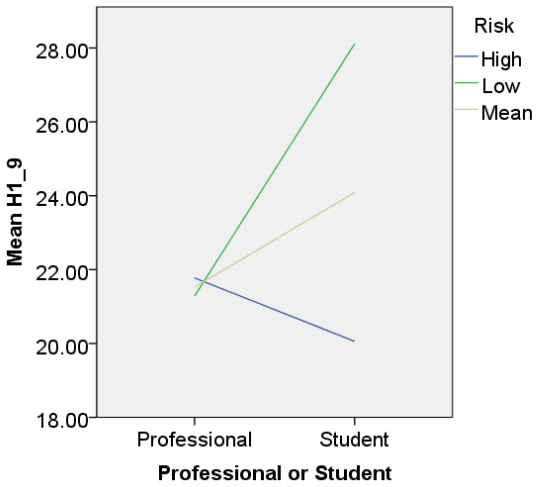

The result that security professionals remain closer to the expected value was also confirmed by the interactions between the variable 'professional or student' and the variable of WTP with the self-reported risk attitude of the individuals. More precisely, moderation analysis revealed a significant interaction with predictor X = *Student* or *Pro*, outcome variables Y = *WTP*, and moderator the Likert-scale self-reported risk attitude M=*General Risk*, interaction $b = 2.06$, 95% CI $[0.15, 3.97]$, $t = 2.14$, $p = 0.034$ (indicatively, interaction with variable $Y = H_1 9$ is shown in Figure 3). *General Risk* is

the survey variable that corresponds to the self-reporting risk question 'How willing are you to take risks in general?' (low values indicate risk averse and high values risk seeking behaviour).
So, amongst risk seeking individuals, being an information security professional had a significant positive relationship with WTP to avoid a lottery; the effect was reversed amongst risk averse individuals, i.e. amongst risk averse individuals, being a professional had a significant negative relationship with WTP. In other words, amongst risk seeking individuals, professionals were the least risk seeking, and amongst risk averse individuals, professionals were the least risk averse. So, in all cases, professionals were closer to risk neutrality than the general population (indicatively, Figure 3).

4.1.1 (A) Between subjects tests

There were overall fifteen WTP-type lotteries, all with negative-only outcomes. For each of these variables, a new variable called *RiskAversionHx_ y* was computed (with $x = 1$, 2 and $y = 1$ to 12; Figure 2). The new variables expresses the distance of the subject's WTP from the expected value (EV) of each lottery. The values are positive if the subject is willing to pay more than the actual expected value, and negative otherwise. So, positive values of this variable imply risk aversion and negative values reveal risk seeking behaviour. A risk neutral subject would have *RiskAversion = 0*.

The test that was used was the parametric one-sample t-test. This test determines whether a sample belongs to a population of a specific mean; the mean in our case was the expected value of the lotteries, but since *RiskAversionHx_ y* variables were computed from the expected values of each group of lotteries, the actual values of the new variables had zero as a reference point. As a result, all t-tests examined mean deviation from zero. The four assumptions for using the one-sample t-test require that the dependent variable is measured at interval or ratio level, data need to be independent, the number of significant outliers needs to be restricted and, lastly, the dependent variable needs to approximate the normal distribution. All assumptions were met since the dependent variable was WTP, measurement was between subjects and sample outliers were shown to approximate the normal distribution (see Appendix B.2) [6].

For the lotteries of *Group A* ($\mu = -2.5$) (Appendix A.1), the one-sample t-test revealed significant risk aversion for all lotteries for both professionals and students. The same result of significant risk aversion was observed for *Group B* ($\mu = -7.5$). However, in *Group C* ($\mu = -25$) statistically significant risk seeking behaviour was observed in one out of the four lotteries of

[6]The t-test is robust against violations of normality, nevertheless, we showed that outliers are distributed roughly normally.

each sample. We can see the positive differences of the mean in Table 3 (risk aversion) and how they become negative from the ninth lottery and on (risk taking). The last three lotteries with large losses did not reveal significant risk attitudes. It is noteworthy that students were willing to pay significantly less than the expected value to avoid lottery $H_1 12$ (ambiguous in both probabilities and outcomes), whereas professionals were significantly risk taking in the simply risky lottery $H_1 9$ (see the list of lotteries in Appendix A.1).

Table 3: One-Sample Test for between-subjects risk aversion

		One-Sample Test (Test Value = 0)					
		Students (df=57)			**Professionals (df=53)**		
	EV	μ diff	95%CI of diff		μ diff	95%CI of diff	
			Lower	Upper		Lower	Upper
$H_1 1$	-2.5	5.22414***	2.7774	7.6709	5.09259***	2.4065	7.7787
$H_1 2$	-2.5	6.62069***	3.7917	9.4496	4.77778***	2.5671	6.9885
$H_1 3$	-2.5	8.17241***	4.9194	11.4255	5.14815***	3.0809	7.2154
$H_1 4$	-2.5	10.25862***	6.6372	13.8800	6.01852***	3.8721	8.1649
$H_1 5$	-7.5	3.77586**	1.5433	6.0085	2.31481*	.0201	4.6096
$H_1 6$	-7.5	6.41379***	4.0169	8.8107	4.66667***	2.1439	7.1894
$H_1 7$	-7.5	5.96552***	3.1745	8.7565	3.94444**	0.0000	6.0928
$H_1 8$	-7.5	8.50000***	5.9158	11.0842	6.11111***	3.9183	8.3040
$H_1 9$	-25	-1.03448	-4.0101	1.9411	-3.48148**	-5.9473	-1.0157
$H_1 10$	-25	-2.22414	-4.9647	.5164	-1.18519	-4.2028	1.8324
$H_1 11$	-25	-1.53448	-4.7128	1.6438	-1.50000	-4.8234	1.8234
$H_1 12$	-25	-3.25862*	-6.2620	-.2552	-.79630	-3.9081	2.3155
$H_2 6$	-86.6	34.08966	-7.8670	76.0463	14.96364	-12.8594	42.7866
$H_2 7$	-86.6	31.03793	-5.3283	67.4042	3.10909	-19.9752	26.1934
$H_2 8$	-89.75	18.85345	-12.5768	50.2837	2.08636	-21.1135	25.2862

* $p \leq 0.05$, ** $p \leq 0.01$, *** $p \leq 0.001$
Test is performed on $H_i j$ variables' WTP differences from each lottery's expected value.

4.1.2 (B) Within subjects tests

The within-subjects design increases the sensitivity of observed effects, as it was the same participants that provided the data for the various conditions. The tests used for these within-subject comparisons were the non-parametric Friedman test [31] that is used to test differences between more than two conditions having a dependent variable of ordinal or continuous type, and the non-parametric Wilcoxon signed rank test [67] that also reveals the magnitude of lottery-pairwise WTP differences. The tests require that the variables are related, i.e. that the same subjects provide the scores for the conditions. The Friedman test ranks all the conditions for each subject separately, and then sums up the ranks for each condition. The independent variables were the expected values of all WTP lottery questions: $H_1 1$ to $H_1 12$. The dependent variable was the amount that individuals are willing-to-pay in order to

avoid each lottery. Lotteries are categorised, based on their expected values into the aforementioned groups.

For *Group A* both non-parametric tests revealed that students have significantly different (increased) WTP amongst the pairs of lotteries, but professionals were more 'robust', i.e. they only showed significantly different behaviour between the risky and the fully ambiguous pair (ambiguous in both probabilities and outcomes; Figure 5), whereas students also revealed significant differences amongst other pairs (Figure 4). The numerical values on the diagram nodes of all lottery pairwise comparisons indicate the sample average rank for each lottery of the group by the Friedman test.

Figure 4: Pairwise comparison of Group A lotteries for Students

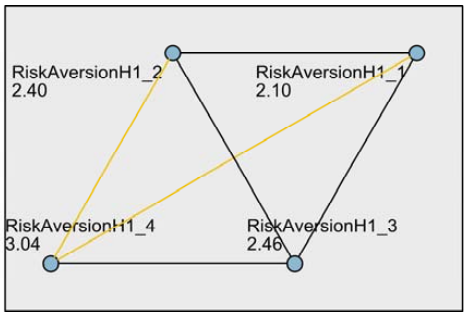

Figure 5: Pairwise comparison of Group A lotteries for Professionals

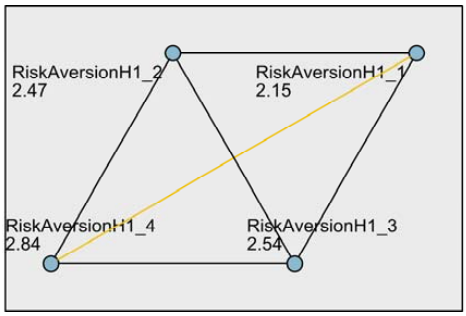

The aforementioned differences between students and professionals of *Group A* were qualitatively balanced in *Group B*, where both samples revealed differences in the same pairs amongst lotteries of different nature (Figures 6 and 7). Moreover, professionals stated a WTP that was very similar for probability- and outcome-ambiguous lotteries.

In *Group C* we observed that students did not diversify their WTP significantly due to ambiguity, but professionals significantly changed their WTP amongst the pairs of lotteries: ($H_1 9$, $H_1 12$) and ($H_1 10$, $H_1 12$) (Figure 8).

Conclusively, we observed that for both samples, as expected, WTP for avoiding a risky lottery was significantly smaller than for avoiding a lottery that was ambiguous in both probabilities and outcomes. It is not clear, however, whether ambiguity of probabilities increased WTP more than ambiguity of outcomes. Professionals were equally or more prone than students to increase their WTP in order to avoid mean preserving spreads of risky lotteries.

Figure 6: Pairwise comparison of Group B lotteries for Students

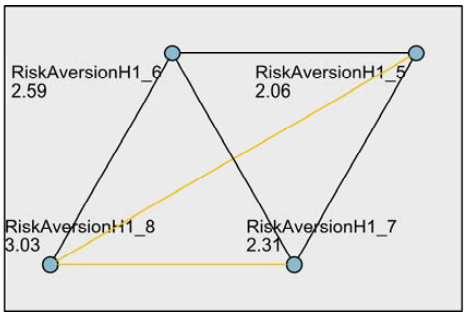

Figure 7: Pairwise comparison of Group B lotteries for Professionals

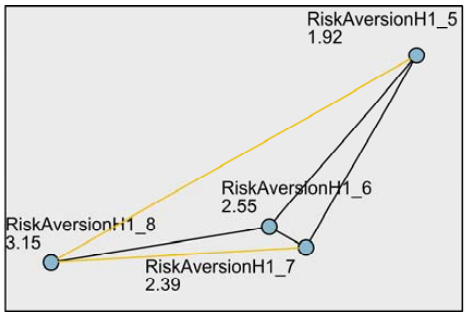

Figure 8: Pairwise comparison of Group C lotteries for Professionals

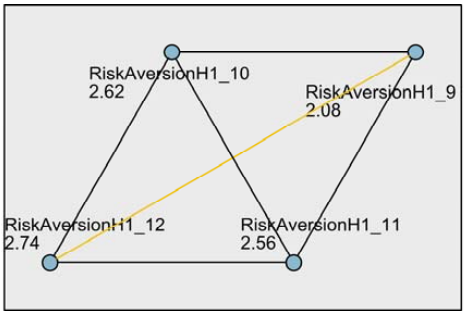

4.2 Worst-case thinking

This section is divided into the analysis of three parts: lottery comparisons, lottery comparisons against stated WTP, and salience theory calculations for each lottery comparison.

4.2.1 Lottery Comparisons and findings on potential heuristics

Finding 4: Both professionals and students reveal choice preferences that are in line with expected values and state-by-state comparisons of lotteries.

Subjects were initially asked to state their WTP in order to avoid three five-outcome lotteries (see Appendix A.3). At a different stage of the experiment subjects were presented with five pairs of lotteries and were asked to chose the one that they prefer (see Appendix A.2). The lotteries of each comparison pair have different attributes, e.g. they vary in their expected value, variance, best (least worse) and worst outcome. Depending on the lottery of each pair that was chosen by each sample, we examined whether

this choice is in accordance, or in contradiction with the relevant attribute. In Table 4 we can see preference percentages per comparison for both samples, as well as the 'fit' of the various heuristics to the given preferences. The major qualitative difference that we observed between professionals and students regarded the first comparison (Lotteries 9 and 10, Appendix A.2). The third comparison was only quantitatively different amongst the two samples. However sample differences were not statistically significant (Pearson's chi-2 test).

We can observe (Table 4) that in the comparisons in which expected value is different for each lottery, the lottery with the smallest expected loss was always chosen. Thus, the possibility that choice is based on the expected value is sustained. If lottery preferences are examined by the variance of the distribution of each lottery, we see that preferences are balanced. That is, choosing the lottery with the smallest variance was not clearly preferred as a heuristic.

Examining the *best possible outcome* of each comparison, i.e. the least damaging loss, we observed that in most of the cases preferences of both professionals and students were in line with this heuristic. One might argue that these choices reinforce expected utility as a heuristic, as three out of the five lotteries approximate power-law distributions, and therefore their smallest losses were associated with large probabilities ($p = 0.85$). However, such distributions underlie the lotteries of the first, second and fifth comparison and it was exactly these lotteries that did not comply with this simple heuristic.

In a similar fashion, the *worst outcome* column examines whether subjects avoided the lottery with the worst-outcome and chose the opposite lottery. It is notable that in all cases except one, the lottery with the largest loss was chosen by both professionals and students. This is arguably not surprising, as this heuristic is actually very simplistic.

Most salient pair is a potential heuristic that was examined under the assumptions of salience theory. There are two separate columns for this choice rule. In the first column we assumed 'same dice roll' and salience was calculated by comparing all pairs of outcomes of *the same state* amongst the two lotteries and specifying the *most salient pair*. The most salient pair was the one that had a larger value of salience function $\sigma(x, y)$ for outcomes x and y (Equation 8). The most salient pair practically means that the difference of the involved outcomes is the most 'noticeable' of all the differences, and consequently the subject chooses the lottery with the smallest loss. Note, that same states correspond to the same probabilities in the compared lotteries. Same 'dice roll' means that if, say, the worst outcome materialises in

Table 4: Lottery comparisons and accordance with heuristics

Lottery pair	Expected Value	Variance	Worst out-come	Best out-come	# of dom-inant states	Most salient pair (same dice roll)	Most salient pair (indep. dice rolls)
	Students						
L_9 VS L_{10} (47%, 53%)	-	✓	×	×	-	✓	×
L_{10} VS L_{11} (60%, 40%)	✓	✓	×	×	✓	✓	×
L_8 VS L_6 (48%, 52%)	-	×	×	✓	-	×	×
L_6 VS L_7 (60%, 40%)	-	×	×	✓	-	×	×
L_4 VS L_{10b} (52%, 48%)	✓	✓	×	✓	✓	✓	×
	Professionals						
L_9 VS L_{10} (58%, 42%)	-	×	✓	✓	-	×	✓
L_{10} VS L_{11} (55%, 45%)	✓	✓	×	×	✓	✓	×
L_8 VS L_6 (33%, 67%)	-	×	×	✓	-	×	×
L_6 VS L_7 (58%, 42%)	-	×	×	✓	-	×	×
L_4 VS L_{10b} (51%, 49%)	✓	✓	×	✓	✓	✓	×

'✓': preference justifies heuristic
'-': heuristic does not influence choice
'×': preference contradicts heuristic predictions
Pairs of percentages indicate preference for each lottery above

the future for lottery A, then the worst outcome will materialise for lottery B too. So, in this heuristic the decision-maker compares the lotteries 'line by line'. It can be argued that presentation of the comparisons (Appendix A.2) encouraged the aforementioned rule of thumb for the decision-makers, as the states of the lotteries under comparison were presented one next to the other. However, results did not sustain such a decision rule, as preferences do not clearly favour the lottery with the smallest loss in the most salient pair. A closer look at the lottery distributions gives some indication that individuals might actually be expected utility maximisers. Comparisons 3 and 4 were never in accordance with the most-salient-pair rule, but the majority of comparisons: first, second and fifth, which follow power-law distributions, were. Since the first states were very probable, choices might imply that the decision-maker not only compares 'line by line', but also sums the outcomes when moving from one line to the next. For example, in the first comparison of Appendix A.2, the decision-maker, when reaching the second line, might add probabilities ($p_1 + p_2 = 0.93$) and since the combination of the first two states gives a very likely event, might choose the cumulatively smallest loss.

Similar reasoning holds for the *most salient pair* on 'independent dice rolls'. This heuristic allows for the two lotteries to happen independently (independent 'dice rolls'), so that in lottery A, say, the best outcome might materialise but in lottery B, say, the worst. The difference here was that the most salient pair was calculated from all possible outcome-combinations amongst the two lotteries. The reasoning behind this particular heuristic was that, by fixing the least-worst outcomes to very similar values, it was actually the worst-case catastrophic outcome that potentially attracted the attention of the decision-maker. We can observe in Table 4 that the majority of results did not favour such a decision rule; there was just some indication that this heuristic complies with some choices of the professionals.

Number of dominant states is the sum of the same-dice-roll states that are strictly preferable to the corresponding states of the opposite lottery. The 'same dice roll' requirement is important here, as it is the corresponding states of 'line-by-line' comparison that produce preference for one of the two lotteries. Note that not all lottery comparisons have a lottery that dominates the opposite lottery in the number of states, as in three of the comparisons lotteries have the same number of dominant states (with either one or three identical states). Only the second and fifth comparisons have a states-dominant lottery. As we can see in Table 4 both these comparisons complied with this heuristic, for both samples. Thus, the the lottery with the most dominant states was preferred by all participants.

4.2.2 Preferred lotteries and stated willingness-to-pay

Finding 5: Security professionals exhibit preference inconsistencies between willingness-to-pay and choice tasks, similarly to the general population.

There was an interesting finding pertaining lottery comparisons and WTP. For the three lotteries involved in the first two comparisons (Lotteries 9, 10 and 11, Appendix A.2) participants also stated their willingess-to-pay to avoid them. We could thus check the consistency of these replies. For the first comparison of L_9 against L_{10}, two variables were created, CONSISTENCY_L_9 and CONSISTENCY_L_{10}vsL_9. In case a subject prefered L_9 to L_{10} in the comparison and was willing to pay less to avoid L_9 than to avoid L_{10}, the subject's replies were consistent and they were coded with a variable value of 0. In case of inconsistency, the value was CONSISTENCY_L_9=1. Similarly, any contradiction regarding L_{10} was examined. So, inconsistency here is the phenomenon of preferring one lottery (from another) and at the same time be willing to spend more to avoid this lottery (than the other). The same reasoning was applied to the comparison and WTP between L_{10} and L_{11}. Note that L_{10} was used in both comparisons, and therefore there were

two variables for L_{10}, one for each comparison. Table 5 depicts the percentage of subjects across both samples that chose L_i over L_j and revealed inconsistency by their stated WTP.

Table 5: Lottery comparisons and willingness-to-pay inconsistencies

Comparison	Preference $L_i \succ L_j$	% of subjects that chose L_i over L_j and revealed choice inconsistency	
		Students	Professionals
L_9 VS L_{10}	$L_9 \succ L_{10}$	33%	47%
	$L_{10} \succ L_9$	58%	52%
L_{10} VS L_{11}	$L_{10} \succ L_{11}$	57%	43%
	$L_{11} \succ L_{10}$	13%	32%

There is no statistically significant difference amongst professionals and students, as inconsistency percentages are high in both samples. So, professionals were by no means found to be more choice-consistent than the general population.

4.2.3 Salience Theory calculations for each lottery-comparison

Finding 6: The majority of security professionals have a distorted perception of probabilities. The general population reveals overall more consistent preferences than security professionals.

Salience theory is a theory of choice among lotteries that quantifies the decision weights of salient lottery outcomes, and proposes that the attention of the decision-maker is focused on the most salient outcomes. Such a focus favours the corresponding salient lottery for positive outcomes and disfavours it when lottery outcomes are in the domain of losses. For the purposes of the analysis of this section, it is assumed that the claims of salience theory [16] are true, and consequently conclusions on the subjects' *local thinking* are derived from the experiment results. Local thinking is defined as the phenomenon in which decision makers do not consider all information that is available to them, but tend to overemphasize the information that their mind focuses on [34].

Methodology for calculating salience theory-predicted preferences over two lotteries can be summarised in the following steps:

- Step 1: write all possible state space pairs by combining all outcomes from the first and the second lottery.

- Step 2: rank all pairs by their salience σ:

$$\sigma(x_s^i, x_s^{-i}) = \frac{|x_s^i, x_s^{-i}|}{|x_s^i| + |x_s^{-i}| + \theta}. \tag{8}$$

A salience function serves as the connecting link between the cognitive notion of salience and a number of properties. These properties are ordering, diminishing sensitivity and reflection and any function that maintains these properties is eligible. The vector containing the payoffs of the lotteries in state s is $x_s = (x_s^i)_{i=1,2}$ and x_s^{-i} is the state s-outcome of lottery L_j, where $j \neq i$. Parameter θ is estimated as $\theta = 0.1$ ([16], page 24) [7].

- Step 3: assign a number k to each pair, starting from the most salient pair. For example, the most salient pair across all states $\sigma(x_s^{max}, x_{s'}^{min})$ has $k = 1$.

- Step 4: compute:

$$\sum_{s \in S} \delta^{k_s} \pi_s [\upsilon(x_s^1) - \upsilon(x_s^2)], \tag{9}$$

where, π_s is the smallest probability of the two outcomes of the pair. Note that the utility function $\upsilon(.)$ has to be linear, for calculating the differences $\upsilon(x^1) - \upsilon(x^2)$.

For example, for two lotteries L_i and L_j, $L_i \succ L_j$ if and only if the sum (9) if positive. An important part of the calculation is the value of $\delta \in (0, 1]$, which expresses the degree of local thinking for a decision-maker. For $\delta = 1$, the decision-maker's probability weighting is exactly the objective probabilities. For $\delta < 1$, local thinking favours the first lottery, L_i, when it 'pays more' in the more salient lottery states. The salient states are the ones that are less discounted by δ due to the exponent k. In our case, only negative outcomes are considered, so $\delta < 1$ favours L_i when it has smaller losses in the most salient states [8].

The following graphs are produced in Mathematica 9.0 [55] and depict the intervals of δ for which the comparison L_i or L_j is expected to reveal preference: $L_i \succ L_j$. Percentages of students and professionals that chose the first lottery and correspond to positive deltas are also given.

[7] Note that outcomes are presented to belong to the same state s here; however, we can assume that of all possible permutations of outcomes, the decision-maker chooses certain ones, so that outcomes of the two lotteries are paired.

[8] It is noteworthy that δ is estimated as $\delta = 0.7$ and that for $\delta = 0.73$ the Allais Paradox is explained by the narrow framing of the local thinker.

Note that here we assume 'independent dice rolls', i.e. pairs are formed by combining all outcomes of the first lottery with all outcomes of the second.

Figure 9: L_9 or L_{10}: values of sum 9 for $L_9 \succ L_{10}$, $\delta \in (0, 1]$
(Students: 47%, Professionals: 58%)

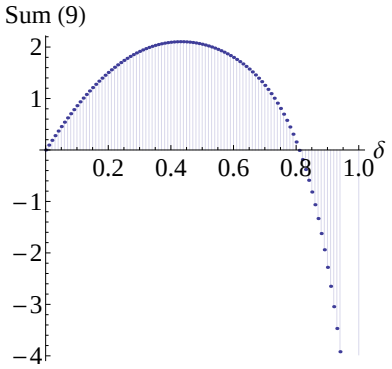

Figure 9 indicates that since the majority of professionals preferred L_9, professionals are associated with $\delta \in (0, 0.8)$. The majority of students chose L_{10}, so they correspond to deltas of negative values, i.e. $\delta \in (0.8, 1]$. This result suggests that decision weights of students are close to objective probabilities (δ close to 1), but preferences of professionals reveal a considerable degree of probability distortion.

Figure 10: L_{10} or L_{11}: values of sum 9 for $L_{10} \succ L_{11}$, $\delta \in (0, 1]$
(Students: 60%, Professionals: 55%)

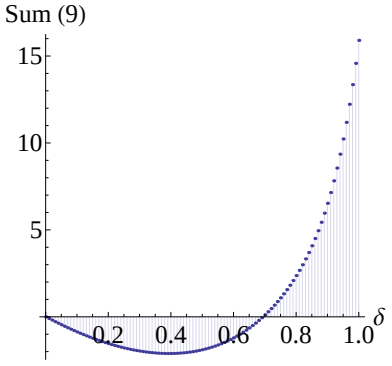

In Figure 10 we see that since the majority of both samples preferred L_{10}, and L_{10} is the first lottery in the comparison (i.e. corresponds to positive δ values), therefore, both professionals and students reveal a $\delta \in (0.7, 1]$, which is an interval that contains objective decision weights.

Figure 11: L_8 or L_6: values of sum 9 for $L_8 \succ L_6$, $\delta \in (0, 1]$
(Students: 48%, Professionals: 33%)

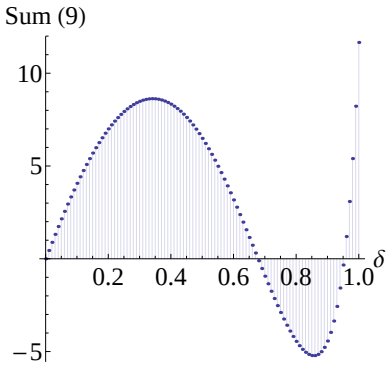

In Figure 11 the majority of both samples prefer the second lottery, therefore, choices correspond to deltas that give negative values, i.e. $\delta \in (0.66, 0.96)$. So, truly objective decision weights are excluded for both students and professionals; their preferences necessarily indicate some probability distortion.

Figure 12: L_6 or L_7: values of sum 9 for $L_6 \succ L_7$, $\delta \in (0, 1]$
(Students: 60%, Professionals: 58%)

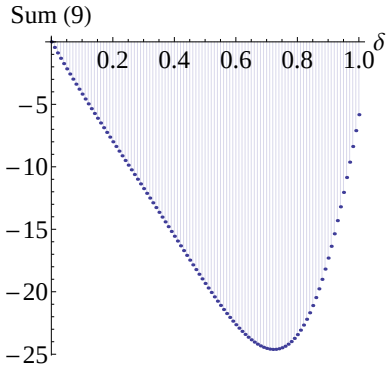

For the fourth comparison (Figure 12), the majority of both samples choose the first lottery, but no additional information is extracted, since the whole range of deltas corresponds to values which have the same sign. So, both objectivity and distortion potentially exist under this choice.

Finally, in Figure 13, the majority of both samples slightly prefer the first lottery, which would give $\delta \in (0.82, 1]$. Similarly to the second lottery comparison (with a more narrow δ interval) this result reveals decision weights even closer to objective probability perception.

Figure 13: L_4 or L_{10b}: values of sum 9 for $L_4 \succ L_{10b}$, $\delta \in (0, 1]$
(Students: 52%, Professionals: 51%)

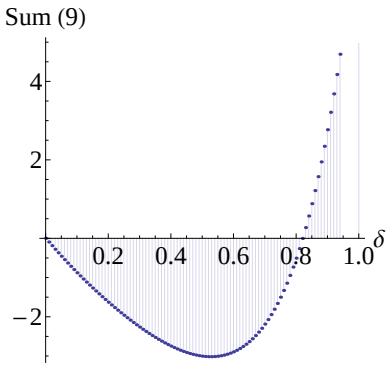

Summarising the results for information security professionals we can see that in the first comparison strong local thinking seems to be prevalent, i.e. $\delta < 0.8$. In the first and the third comparisons objective deltas are completely excluded. Only the remaining three comparisons reveal local thinking that corresponds to delta-intervals which include the value $\delta = 1$, i.e. might imply objective perception of probabilities. However, distance from objective weighting is not negligible: the lowest potential value is approximately $\delta = 0.67$ and in the same comparison (L_8 or L_6) objective weighting of probabilities is excluded, allowing only for $\delta < 0.94$. For the student sample the intersection of the δ intervals is $(0.82, 0.96)$. This means that there is some local thinking, i.e. a distortion of objective probabilities that favours the lotteries that contain smaller losses in salient pairs. Interestingly, and due to the diversified first lottery comparison, intersection of the δ-intervals for professionals is the empty set. This means that preferences of the majority of professionals are not consistent enough to allow for a clear estimation of the degree of their local thinking.

4.3 Other-evaluation Ambiguity Aversion

Finding 7: There is no evidence that subjects change their risk behaviour when they are informed that they will be evaluated by other parties.

No significant differences were observed between the control and the treatment groups of the hypothesis, in either lottery comparisons or WTP questions. The most probable explanation is that it is hard to create a sense of 'evaluation by other parties' in an online environment. That is, participants already knew that their responses were subjected to 'evaluation' for either statistical analysis or validity checks.

4.4 Security - Operability trade-off

Findings on preferences between Security and Operability

Finding 8: Security professionals reveal preferences that favour operability over security. These preferences are significantly dependent on their job role.

When asked to choose between two mechanisms that either enhanced the security of a system or its operational time (with the same monetary values assigned to each of the two attributes), the majority of professionals (58%) preferred operability over security enhancement. However, preferences might have been influenced by the actual information security roles of the professionals, giving them a certain point of view. For that reason, we examined how this preference varies amongst the various job roles and indeed significant diversification between security and operability preference was found across the various positions. The question presented to the participants is included in Appendix A.4 and preferences are shown in Table 6.

Table 6: Security VS Operability preference across Security Job Titles

	Job Title				
	Senior executive role (e.g. CEO, CIO, CISO, CSO etc.)	**Managerial role** (e.g. Project Manager, IT Director, Security Manager etc.)	**IT & Security role** (e.g. Security Officer, System Administrator, Cyber Security Information Analyst etc.)	**Compliance, Risk or Privacy role** (e.g. Governance, Risk & Compliance Consultant, Information Security Consultant, Auditor etc.)	**Other**
Mechanism A Enhances Security of the system by 10% (chosen by 42%)	5	3	7	8	0
Mechanism B Enhances Operability of the system by 10% (chosen by 58%)	1	13	7	3	2

$\chi^2(4, N = 55) = 12.092, p = .017$

Results in Table 6 show that compliance and risk professionals are security-oriented, as might have been expected, due to the certification and regulatory issues that they are exposed to. Also, not surprisingly, professionals with managerial roles preferred operability, as their positions are more project and task-oriented. However, IT professionals expressed a balanced preference between operability and security. Finally, senior executives chose security.

Findings on switching points and loss aversion in Security and Operability

Finding 9: Security-focused professionals insist more on the importance of security over operability than operability-focused professionals.
Finding 10: Professionals that prefer operability to security have a more balanced perception of the two attributes.

This part of the analysis considered the estimation of a switching point between security and operability and the measurement of the magnitude of loss aversion in both security and operability.

Each participant revealed a 'switching point' between security and operability. If the subject initially preferred security to operability, then their consecutive preferences were stored in variable SWITCHPOINT_SEC. In Figure 14, value x denotes a switching point of enhancing security by $x\%$ ($x < 10$) and operability by 10%, after which operability enhancement became more attractive to the subject. So, x can be considered as a 'balance point' for which the utility of $x\%$ of security equals the utility of 10% of operability: $Sec(x\%) = Ops(10\%)$. Figure 15 is the operability equivalent. More precisely, both security-oriented professionals and professionals that chose operability revealed switching points close to the mean, which suggests that they both weighted their favourite attribute 'twice as much' as the attribute they did not choose (Figures 14 and 15). Practically, we could state that, on average, an enhancement of their favourite attribute by $x\%$ has the same utility as an enhancement of the not preferred attribute by $2x\%$.

The second measurement that was performed in this series of questions was the relative loss aversion between security and operability, as described in the design of Hypothesis 4 (section 3.2.4). Variables LOSS_AV_SEC (Figure 16) and LOSS_AV_OPS (Figure 17) measure the difference between the aforementioned switching point and elicited preferences of Scenario 2, which included reduction of the level of one of the attributes. The logic behind this measurement of relative loss aversion amongst the two attributes is the following: sometimes loss, or marginal reduction of an attribute level in our case, 'hurts more' the individuals than an equivalent enhancement 'satisfies'.

Findings suggest that subjects who had a preference for security exhibited *relative loss aversion* between the two attributes (security and operability) and *loss aversion* in the security attribute. More specifically, security-focused professionals weighted reduction of security almost as much as they valued triple the enhancement of operability; this is because the mean of loss aversion in security was $\mu_{LOSS\_AV\_SEC} = 2.15$ and the mean switching point for security was $\mu_{SWITCHPOINT\_SEC} = 5.26$, thus,

Figure 14: Security switching points $(Sec(x\%), Ops(10\%))$

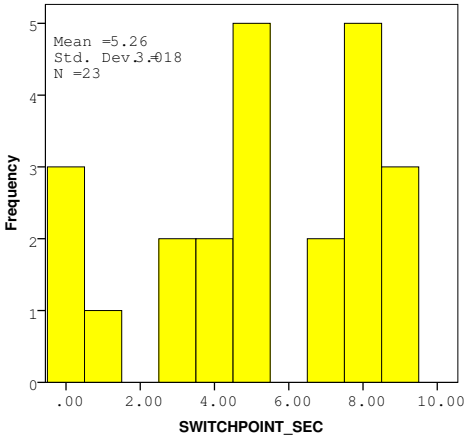

Figure 15: Operability switching points $(Sec(10\%), Ops(x\%))$

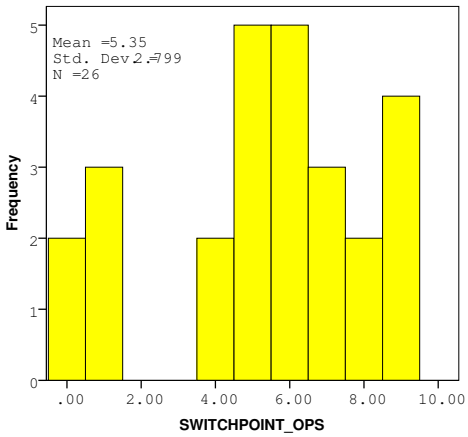

$-Sec(-(5.26-2.15)\%) = -Sec(-3.11\%) > Ops(10\%)$, which means that reduction of security 'hurt' about three times more than enhancement of operability 'satisfied' the decision-makers.
This preference also holds for security: reduction of security was valued almost twice as security enhancement, and since, on average, $Sec(5.26\%) = Ops(10\%)$, thus $-Sec(-3.11\%) > Sec(5.26\%)$. So, reduction of security 'hurt' almost double as its enhancement 'satisfied'. This result is in accordance with prospect theory's loss aversion findings on lotteries with gains and losses.
On the other hand, professionals who chose operability revealed, on average, smaller relative loss aversion between operability and security, as
$\mu_{LOSS\_AV\_OPS} = 1.83$ and $\mu_{SWITCHPOINT\_OPS} = 5.35$, thus,
$-Ops(-(5.35-1.83)\%) = -Ops(-3.52\%) > Sec(10\%)$. Similarly, their loss aversion in operability itself is, on average, much less than double, as
$-Ops(-3.52\%) > Ops(5.35\%)$. Therefore, overall, operational time was preferred to security, but its relative loss aversion to security was much smaller than the equivalent loss aversion that security-focused professionals revealed. So, professionals with a focus on security insisted more strongly on and reacted more to security losses.

Finally, it is noteworthy that professionals who have a preference for operability were more likely to exhibit linear preferences between reduction and enhancement of the attributes in their consecutive choices, as many of them revealed zero loss aversion in operability and it was the mean that produced the final loss aversion result.

Figure 16: Loss Aversion in Security

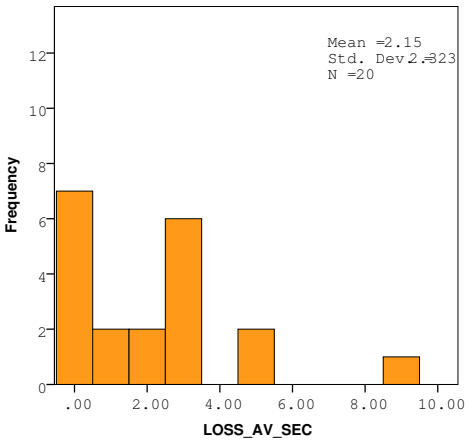

Figure 17: Loss Aversion in Operability

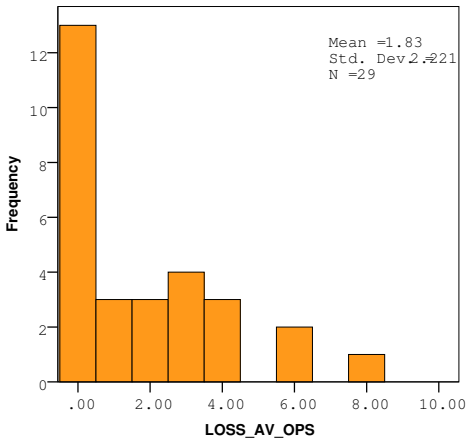

4.5 Survey Analysis

Finding 11: Security professional reveal different risk attitudes to the ones they self-report.
Finding 12: The educational level of participants influences risk attitude on willingess-to-pay.

Table 7: Spearman's correlation coefficients for General Risk

		Students N=58	Pros N=54
$H_1 1$	rho	-.061	.088
	Sig. (2-tailed)	.648	.526
$H_1 2$	rho	-.105	.176
	Sig. (2-tailed)	.431	.203
$H_1 3$	rho	-.053	.186
	Sig. (2-tailed)	.695	.177
$H_1 4$	rho	-.032	.316*
	Sig. (2-tailed)	.814	.020
$H_1 5$	rho	-.056	.160
	Sig. (2-tailed)	.677	.248
$H_1 6$	rho	-.136	.266
	Sig. (2-tailed)	.310	.052
$H_1 7$	rho	-.137	.280*
	Sig. (2-tailed)	.306	.040
$H_1 8$	rho	-.131	.336*
	Sig. (2-tailed)	.327	.013
$H_1 9$	rho	-.294*	-.029
	Sig. (2-tailed)	.025	.837
$H_1 10$	rho	-.325*	.109
	Sig. (2-tailed)	.013	.435
$H_1 11$	rho	-.253	-.077
	Sig. (2-tailed)	.055	.579
$H_1 12$	rho	-.270*	-.024
	Sig. (2-tailed)	.041	.866

A number of analyses were conducted on the survey data and their relations with the experiment results. Some of the findings are presented here. Significant correlation was found between *general risk attitude* and *WTP* for three out of twelve lotteries, for both samples. *General Risk* represents the survey question: 'How willing are you to take risks in general?' (low values indicate risk averse and high values risk seeking behaviour). Student behaviour confirms literature findings on correlation of self-reported risk attitude and actual behaviour [25], but responses of professionals contradict the expected results. We observe (Table 7) that both students and professionals revealed some significant correlations between self-stated risk attitude and WTP. Students behaved as expected, i.e. by revealing negative correlation (significant negative correlation in 3 out of the 12 lotteries), whereas professionals positive (significant positive correlation in 3 out of the 12 lotteries). This implies that, in some cases, professionals who reported themselves as risk taking were actually willing to pay more in order to avoid the lotteries, so they were actually risk averse. This inconsistency was not observed for students; on the contrary, in some cases students' statements were confirmed by their behaviour.

Figure 18: Interaction of *Pro or Student* and $H_1 6$ with *number of family dependents* as moderator

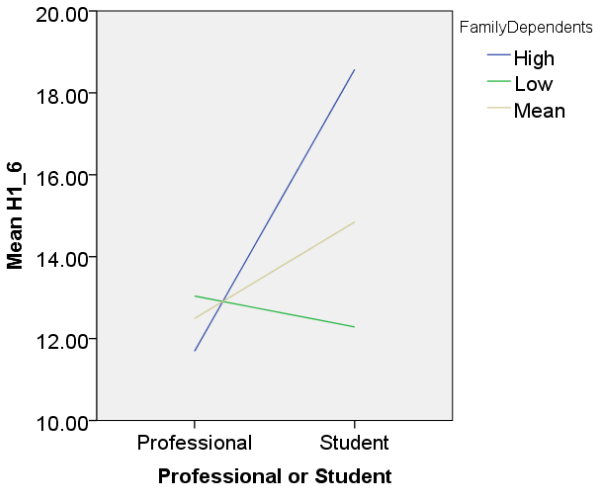

A number of linear regression models were conducted for the analysis of survey and experiment data, but results did not reveal significant predictors. The specifications for the models are described in Appendix C.1.

The demographic variable of the number of family dependents was found to cause an interaction. In particular, moderation analysis revealed a significant interaction between predictor $X = Student$ or Pro and the outcome vari-

ables $Y = WTP$ (indicatively, variable $H_1 6$) and moderator M=*number of Family Dependents*, interaction $b = -3.22$, 95% CI $[-5.92, -0.5]$, $t = -2.37$, $p = 0.019$. In other words, when the number of family dependents was high, being an information security professional had a significant negative relationship with WTP; the effect was observed across all lottery level stakes, except for very high (indicatively, Figure 18). The expected result would be a positive relationship between number of family dependents and WTP, i.e. risk aversion, which was manifested for students, but, surprisingly, did not hold for professionals.

Table 8: Kruskal-Wallis Test with dependent variable WTP and 4 Educational levels

Kruskal-Wallis Test (N=109, df=3)	
Lottery	Test statistic
$H_1 1$	17.809***
$H_1 2$	8.228*
$H_1 3$	8.090*
$H_1 4$	11.243**
$H_1 5$	8.908*
$H_1 6$	3.376
$H_1 7$	6.429
$H_1 8$	8.138*
$H_1 9$	0.943
$H_1 10$	6.996
$H_1 11$	3.765
$H_1 12$	7.611*
* $p \leq 0.05$, ** $p \leq 0.01$, *** $p \leq 0.001$	

The educational level was also found to have a significant effect on WTP. The non-parametric Kruskal-Wallis test on the merged sample reveals significant differences in WTP amongst the four levels of education: highschool, bachelor's degree, master's degree and PhD (Table 8). The overall trend was a higher WTP for participant with bachelors, and significant differences amongst the pairs of highschool-bachelor's and highschool-PhD. This might imply that the observation was caused by the student sample, the subjects of which are most likely at bachelor's level. However, this explanation was rejected as there was no interaction between educational level and attribute 'professional or student' on WTP.

5 Discussion

The scope of this study was to specify behavioural aspects of decision-making under risk and ambiguity that information security professionals exhibit, and to contrast these attitudes against the general population. In other words, we intended to examine whether security professionals are rational decision-makers, and investigate whether certain underlying characteristics of information security shape a unique context. We divided the experiment into

four major hypotheses, containing a number of sub-hypotheses and tests.

Security professionals exhibited significant risk aversion for small losses. This result, for the case of professionals, might mean that they consider small losses inevitable and therefore are willing to pay more to avoid them. This might have implications in a security environment, as such behaviour would always justify measures against low-impact threats. However, these losses are also associated with small probabilities, which could imply that professionals do not want to take risks, even if an event has very little likelihood of materialising.

The observed behavioural pattern of professionals complies with the *four-fold pattern* of risk attitudes for the domain of losses, introduced by Kahneman and Tversky [43]. Based on this pattern, professionals switched from being risk averse and became risk seeking for large probabilities. This finding implies that professionals 'hope' to avoid a very likely loss, and they might consequently reject a favourable settlement. The settlement in this case could be a security investment amount that is equal to the expected loss, which the professionals might refuse to accept, as they would behave in a risk taking manner.

The combination of risk aversion for small-losses and the four-fold pattern could imply that preventive measures for common information security threats (e.g. malware, viruses) are viewed as necessary, unavoidable investment; but it would be quite alarming if professionals were to maintain their risk taking attitude for highly possible threat events. As we argued in the introduction, there is capacity for individual risk attitude to be manifested in the currently accepted risk assessment methodologies.

In all relevant lotteries, professionals were always more alarmed than students when they confronted ambiguous probabilities and outcomes. They expressed this fact by becoming significantly more risk averse. However, it is reassuring that professionals consistently stated WTP closer to the expected losses than students did. Moreover, professionals did not seem to separate between ambiguity in probabilities and ambiguity in outcomes. These findings might indicate a 'robustness' of professionals' against ambiguity. The fact that professionals were alarmed by mean-preserving spreads, but they always remained closer to expected losses, might reflect their familiarity with similar presentation formatting of probabilities and losses.

Analysis on heuristics revealed that expected value and a line-by-line comparison of lotteries are consistent with professionals' choices. All subjects chose the lottery with the number of most dominant states to its counterpart lottery. It cannot be inferred whether subjects used a more complex

rule here, such as an estimation of 'how strong' dominance was in each state. This finding is interesting, because if it holds in general it would imply that decisions could be 'nudged' towards some direction. For example, an even amount of states might promote indecisions, as it would make it easier to have the same amount of dominant states. Another possibility would be to choose the states that represent the distribution of each lottery in such a fashion that favours the choice of one of the two lotteries.
There were only indications that worst-case outcomes influenced the professionals' decisions, so, in this case, it seems that the rule that professionals followed approximates expected utility maximisation.

However, we would not characterise security professionals as rational decision-makers. The inconsistencies they revealed between WTP and lottery comparison tasks were in some cases more contradicting than students' replies. The observed probability distortion, measured by the decision weights that are disproportionately assigned to salient outcomes, was even more puzzling, as the majority of professionals did not even manifest a consistent pattern in the way that students did. So, security professionals are very likely to have a biased perception of probabilities and, moreover, this perception is heavily influenced by the framing or presentation of the problem at hand. This fact implies that calculations involved in risk assessment methodologies are indeed susceptible to the subjective perception of the security decision maker. Thus, this can be considered as another weak link in the security chain that needs to be strengthened. A descriptive pluralism for relevant risk methodologies might be a starting point towards this direction.

Professionals that preferred security over operability were much more adamant in their choice. Operability-focused individuals, however, revealed a more balanced understanding between security and operational time. This could suggest that a portion of the operations-oriented professionals are actually more objective in balancing losses and gains (reduction and enhancement) than their security-focused colleagues. In conjunction with the aforementioned finding on the influence of job position, this fact might imply a relation of operability with a 'more practical' business-oriented approach that allows for a more objective (symmetric) contrasting of gains and losses. Preference of the majority of professionals for operability might again be related with a business-oriented point of view, whereas the focus on security might indicate a more traditional 'IT and Security' approach.
Senior positions are usually associated with risk ownership and liability; also, positions that are higher in the hierarchy are able to see 'the big picture' of the security environment. The fact that these individuals chose security over operability might indicate that professionals of such positions are inclined to consider the potential catastrophic and disastrous outcomes that can disrupt business functions, and therefore choose the 'safer path' of security prioriti-

sation.
In any case, the information security context seems to have a significant role in professionals' behaviour, as well as the risk attitude of the individuals.

6 Conclusion

We conducted an online experiment and a survey in order to test four major hypotheses for understanding risk behaviour of security professionals. Our findings suggest that security professionals have distinctive behavioural characteristics regarding risk.

Professionals are more worried about ambiguity than the general population and are better at estimating losses. Their elicited shifts of attitude from risk averse to risk seeking when losses became more probable would raise concerns in an information security investment setting.

Biases and inconsistencies are also vividly manifested in professionals' behaviour. These inclinations of security decision-makers might cast doubts on the appropriateness of current risk assessment methodologies and their corresponding risk management approaches. The reason is that elements of such methodologies are susceptible to being shaped by individual risk perception.

Finally, the actual security environment is itself a key factor in the manifestation of risk attitudes that should be given greater attention. For that reason, our future research will focus on the examination of risk behaviour in a real-world security context, in order to examine environmental parameters that influence security decisions and to contrast their impact against individual traits.

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] Alessandro Acquisti and Jens Grossklags. What can behavioral economics teach us about privacy. *Digital privacy*, page 329, 2007.

[5] Santosh Anagol, Sheree Bennett, Gharad Bryan, Tiffany Davenport, Nancy Hite, Dean Karlan, Paul Lagunes, and Margaret McConnell. There's something about ambiguity. Working Paper, Yale, 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 Papers of Kenneth J. Arrow)*, volume 4. Cambridge, Massachusetts: Belknap 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] Adrian Baldwin, Yolanta Beres, Geoffrey B Duggan, Marco Casassa Mont, Hilary Johnson, Chris Middup, and Simon Shiu. Economic methods and decision making by security professionals. In *Economics of Information Security and Privacy III*, pages 213–238. Springer, 2013.

[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] Aaron Clauset, Cosma Rohilla Shalizi, and Mark EJ Newman. Power-law distributions in empirical data. *SIAM review*, 51(4):661–703, 2009.

[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] Jim Engle-Warnick, Javier Escobal, and Sonia Laszlo. Ambiguity aversion as a predictor of technology choice: Experimental evidence from Peru. *CIRANO-Scientific Publications 2007s-01*, 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] Dinei Florêncio and Cormac Herley. Sex, lies and cyber-crime surveys. In *Economics of Information Security and Privacy III*, pages 35–53. Springer, 2013.

[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)*, 5(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] William T Harbaugh, Kate Krause, and Lise Vesterlund. The fourfold pattern of risk attitudes in choice and pricing tasks*. *The Economic Journal*, 120(545):595–611, 2010.

[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] Provo Qualtrics. Qualtrics software, Version 37,892. Provo, Utah, USA., 2013.

[55] Wolfram Research. Inc., Mathematica, Version 9.0, Champaign, Illinois, USA, 2012.

[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.

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

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

[62] 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.

[63] Shelley E. Taylor and Suzanne C. Thompson. Stalking the elusive "vividness" effect. *Psychological Review*, 89(2):155, 1982.

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

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

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

[67] Frank Wilcoxon, SK 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.

A Appendix - Experiment Design

A.1 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?'.

'What is the maximum amount that you are willing to pay in order to avoid playing a lottery in which there is a probability between p_1% and p_2% of losing \$50?'.

'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 an amount between \$$x_1$ and \$$x_2$ and losing nothing otherwise?'.

'What is the maximum amount that you are willing to pay in order to avoid playing a lottery in which there is 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	**Exp. Outcome Interval**	**Outcome Range**
$H_1 1$	5	-50	0 to 100	-2.5	-2.5	0
$H_1 2$	0-10	-50	0 to 100	-2.5	[-5, 0]	5
$H_1 3$	5	-80 to -20	0 to 100	-2.5	[-4, -1]	3
$H_1 4$	0-10	-80 to -20	0 to 100	-2.5	[-8, 0]	8
$H_1 5$	15	-50	0 to 100	-7.5	-7.5	0
$H_1 6$	0-30	-50	0 to 100	-7.5	[-7.5, 0]	7.5
$H_1 7$	15	-80 to -20	0 to 100	-7.5	[-12, -3]	9
$H_1 8$	0-30	-80 to -20	0 to 100	-7.5	[-24, 0]	18
$H_1 9$	50	-50	0 to 100	-25	-25	0
$H_1 10$	35-65	-50	0 to 100	-25	[-32.5,-17.5]	15
$H_1 11$	50	-80 to -20	0 to 100	-25	[-40, -10]	30
$H_1 12$	35-65	-80 to -20	0 to 100	-25	[-52, -7]	45

A.2 Lottery Comparisons

Hypothesis 2 Question 1 ($H_2 1$)	
Lottery A (Lottery 9)	**Lottery B** (Lottery 10)
a probability of 85% of losing 45	a probability of 85% of losing 50
a probability of 8% of losing 220	a probability of 8% of losing 170
a probability of 3.5% of losing 300	a probability of 3.5% of losing 300
a probability of 2.5% of losing 450	a probability of 2.5% of losing 400
a probability of 1% of losing 900	a probability of 1% of losing 1000
μ = -86.6, Var = 14406.2	μ = -86.6, Var = 14087.4

Hypothesis 2 Question 2 ($H_2 2$)	
Lottery A (Lottery 10)	**Lottery B** (Lottery 11)
a probability of 85% of losing 50	a probability of 85% of losing 45
a probability of 8% of losing 170	a probability of 8% of losing 250
a probability of 3.5% of losing 300	a probability of 3.5% of losing 350
a probability of 2.5% of losing 400	a probability of 2.5% of losing 450
a probability of 1% of losing 1000	a probability of 1% of losing 800
μ = -86.6, Var = 14087.4	μ = -89.75, Var = 14416.2

Hypothesis 2 Question 3 ($H_2 3$)	
Lottery A (Lottery 8)	**Lottery B** (Lottery 6)
a probability of 15% of losing nothing	a probability of 15% of losing nothing
a probability of 30% of losing 200	a probability of 30% of losing 166.66
a probability of 30% of losing 300	a probability of 30% of losing 300
a probability of 20% of losing 450	a probability of 20% of losing 450
a probability of 5% of losing 700	a probability of 5% of losing 900
μ = -275, Var = 28375	μ = -274.998, Var = 40708.8

Hypothesis 2 Question 4 ($H_2 4$)	
Lottery A (Lottery 6)	**Lottery B** (Lottery 7)
a probability of 15% of losing nothing	a probability of 15% of losing nothing
a probability of 30% of losing 166.66	a probability of 30% of losing 183.33
a probability of 30% of losing 300	a probability of 30% of losing 300
a probability of 20% of losing 450	a probability of 20% of losing 450
a probability of 5% of losing 900	a probability of 5% of losing 800
μ = -274.998, Var = 40708.8	μ = -274.999, Var = 33958.5

Hypothesis 2 Question 5 ($H_2 5$)	
Lottery A (Lottery 4)	**Lottery B** (Lottery 10_b)
a probability of 85% of 50	a probability of 85% of 46
a probability of 8% of losing 150	a probability of 8% of losing 180
a probability of 3.5% of losing 300	a probability of 3.5% of losing 350
a probability of 2.5% of losing 450	a probability of 2.5% of losing 480
a probability of 1% of losing 1000	a probability of 1% of losing 900
μ = -86.25, Var = 14698.4	μ = -86.75, Var = 15012.5

A.3 H2 Willingness-to-pay Lotteries

Hypothesis 2 Question 6 ($H_2$6)
Lottery 9: How much are you willing to pay in order to avoid playing a lottery in which there is:
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

Hypothesis 2 Question 7 ($H_2$7)
Lottery 10: How much are you willing to pay in order to avoid playing a lottery in which there is:
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 = 14087.2

Hypothesis 2 Question 8 ($H_2$8)
Lottery 11: How much are you willing to pay in order to avoid playing a lottery in which there is:
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

A.4 Survey Questions

- Question: 'Are you related with the profession or practice of Information Security in any way?' *Yes / No*
- Question: 'How many years of experience do you have in Information Security related tasks?'
- Question: 'How willing are you to take risks in general?' *0 to 10*
  0: Not willing at all 10: Very willing
- Question: '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*

- Question: 'Does your job position allow you to make independent Information Security related decisions?' *Yes / No*

- Question: '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

- Question: 'How worried are you about new unidentified information security threats?' *0 to 10*
  0: Not worried at all 10: Very worried

- Question: 'Have you experienced any important security incident in the past?' *Yes / No*

- Question: '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

- Question: '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

- Question: 'Where / to whom does your Chief Information Security Officer (CISO or CSO) or equivalent senior executive report?'

- Question: 'What is the size of your company?'

- Question: 'What is your gender?'

- Question: 'What is your age?'

- Question: 'What is your educational level?'

- Question: 'What is your marital status?'

- Question: 'What is the number of dependents in your family?'

- Question: 'What is your approximate annual income in British pounds?'

- Question: 'Which country do you live in?'

- Question: 'What is your nationality?'

- Question: 'What is your mother tongue?'

B Appendix - Experiment Analysis

B.1 Data Cleaning

Data analysis was conducted using SPSS version 21 [1] and data cleaning consisted of the following actions:

1. There were two datasets collected for the purposes of this experiment. The first dataset was collected between 21/05 and 11/06/2014 and it was targeted at alumni and MSc students at Royal Holloway. The majority of the participants are information security professionals. The second sample was collected on 26/08/2014 and was targeted at the student database of the Laboratory for Decision Making and Economic Research at Royal Holloway, University of London. The majority of this sample consisted of individuals that are not related to information security. Datasets were combined.

2. A filter was implemented by the use of the willingness-to-pay (WTP) questions of Table 9. Half questions of the table have a maximum monetary loss of 50 USD and the other half a maximum loss of 80 USD. Replies with values greater than fifty and eighty dollars respectively, have been excluded from the analysis of the corresponding lotteries. Only a few cases were excluded from the analysis by using this filter, by being considered invalid; in all these cases, there were consecutive willingness-to-pay choices to avoid lotteries that were larger than the maximum potential loss.

3. All missing cases were excluded. These were caused either by subjects that aborted the experiments half-way or subjects that happened to be online when the experiment became inactive.

The final valid number of cases was $N_1 = 55$ for professionals, $N_2 = 58$ for students, and $N = 113$ for the merged dataset.

An additional validity check was conducted on the significance of the variable *mother tongue*, to see whether non-native English speakers had any issues with understanding instructions or questions. No language effect was found in the data.

B.2 Outliers

For testing whether there is a significant number of outliers in the sample, we used the following method. The z-scores were computed for all WTP questions of variables $H_1 i$ and $H_2 j$. Then the cumulative percentage of cases that had a standard deviation that was larger in absolute value than 1.96 was computed. If this percentage constituted more than 0.05 of the total cases, then there would be more outliers in the distribution of the given variable than we would expect in a normal distribution. It was however important that this analysis was conducted separately for professionals and students, so that we can exclude the possibility of having the sample type act as a moderator; for this reason the merged dataset was split into two. We should state that no outliers were excluded by this methodology, the purpose of which was to examine their distribution.

The analysis revealed six out of the fifteen variables ($H_1 1, H_1 2, H_1 3$, $H_1 4$, $H_1 7$ and $H_2 8$) with outlier percentages more than the expected. However, at closer examination we observed that this deviation was caused by one or two large values in the whole sample. Moreover, the aforementioned variables either had only one or no extreme values ($|z| > 3.29$) and the majority of potential outliers was in the range of $|z| \in (1.96, 2.58)$ or $|z| \in (2.58, 3.29)$. Therefore, the existence and distribution of outliers can be considered roughly within the expected ranges of a normal distribution. This means that existence of outliers was at the edge of being considered significant, and the following statistical tests on the data could be conducted without considering additional 'without-outlier' analyses.

It is also worth noting that the deviation from normality by outlier values was mainly observed in the lotteries with low expected value where higher WTP values could occur more easily.
Table 10 contains the percentages of the values that are potential outliers for all outcome variables, split into students and professionals. Cumulative percent denotes the exact portion of data cases that have z-scores, such that $|z| > 1.96$. Valid percent is the portion of cases in the range $1.96 < |z| < 2.58$. So, a difference between valid and cumulative percentage implies the existence of more extreme outliers, i.e. with z-scores $|z| > 2.58$.

Table 10: Potential Outliers ($|z| > 1.96$) for the z-scores of all outcome variables

	Students		Professionals	
Variable	Valid Percent	Cumulative Percent	Valid Percent	Cumulative Percent
$H_1 1$	1.7	6.9	3.7	3.7
$H_1 2$	3.4	8.6	3.7	5.6
$H_1 3$	3.4	6.9	5.6	7.4
$H_1 4$	6.9	10.3	3.7	7.4
$H_1 5$	3.4	6.9	3.7	3.7
$H_1 6$	1.7	3.4	3.7	3.7
$H_1 7$	3.4	5.2	3.7	5.6
$H_1 8$	3.4	3.4	1.9	3.7
$H_1 9$	3.4	3.4	1.9	5.6
$H_1 10$	3.4	5.2	3.7	3.7
$H_1 11$	3.4	3.4	5.6	5.6
$H_1 12$	3.4	3.4	3.7	3.7
$H_2 6$	1.7	5.2	3.6	3.6
$H_2 7$	1.7	5.2	1.8	3.6
$H_2 8$	5.2	8.6	3.6	5.5

B.3 Controlling for Order Effects

Before measuring the actual attitudes on risky and ambiguous lotteries, we examined data for potential order effects. In order to control for potential order effects in the series of $H1_i$ instrument variables, two conditions were created in the experiment, one presenting the risky lotteries first and then progressing to the ambiguous lotteries and another condition with the opposite order.

Subjects were randomly assigned to one of these two conditions. The first group was named *Risk-to-Ambiguity* group, was marked with a dummy variable $RISK\_FIRST = 1$, and contained questions $H_1 1$, $H_1 5$, $H_1 9$, $H_1 3$, $H_1 7$, $H_1 11$. The second group, the *Ambiguity-to-Risk* one, consisted of lottery-questions $H_1 4$, $H_1 8$, $H_1 12$, $H_1 2$, $H_1 6$, $H_1 10$.
Since, there are two conditions with different subjects, analysis on these two groups was conducted by the non-parametric Mann-Whitney test, and the sample was split into professionals and students, using a filter variable that asks participants whether they are related to the Information Security profession.

Both professionals and students samples were found free of any order effect between risk and ambiguity, as there was no statistically significant difference between the two condition groups (Table 11).

Table 11: Mann-Whitney U Test for Order Effects

		Students N=58	Professionals N=54
$H_1 1$	Test Statistic	316.5	294
	Sig. (2-tailed)	.101	.219
$H_1 2$	Test Statistic	307	259
	Sig. (2-tailed)	.075	.064
$H_1 3$	Test Statistic	395.5	347.5
	Sig. (2-tailed)	.701	.767
$H_1 4$	Test Statistic	411.5	318.5
	Sig. (2-tailed)	.894	.422
$H_1 5$	Test Statistic	371.5	332.5
	Sig. (2-tailed)	.449	.576
$H_1 6$	Test Statistic	349.5	292
	Sig. (2-tailed)	.270	.207
$H_1 7$	Test Statistic	475	367.5
	Sig. (2-tailed)	.390	.958
$H_1 8$	Test Statistic	411	299
	Sig. (2-tailed)	.888	.255
$H_1 9$	Test Statistic	468.5	279.5
	Sig. (2-tailed)	.444	.132
$H_1 10$	Test Statistic	442	346
	Sig. (2-tailed)	.730	.748
$H_1 11$	Test Statistic	449	280
	Sig. (2-tailed)	.649	.143
$H_1 12$	Test Statistic	483.5	310
	Sig. (2-tailed)	.321	.344

Is distribution of $H_1 i$ the same across categories of
'Risky questions presented before Ambiguity questions'?
Null hypothesis is retained for all variables, for both samples.

C Appendix - Regression Specifications

C.1 Linear Models Regression Specifications

We conducted a number of regressions with bootstrapping on all survey variables, by the following specifications. In the initial three regression models the dependent variable is willingness-to-pay (WTP), i.e. the series $H_1 i$ and variables $H_2 6, 7$ and 8.

Specification 1: explores potential differences between the population of professionals and the general population (students). The predictors used in the model are the clearly exogenous variables.
Dependent variable: all variables of Table 9 (H1 Instrument) and variables $H_2 6, 7$ and 8.
Predictors: age, gender, education, marital status, number of dependents in family, country, nationality, language.
Sample: professionals and students.

Specification 2 is the same as Specification 1, having only the additional variable of general risk ('How willing are you to take risks in general?).

Specification 3 aims to explore potential differences amongst the population of professionals. The predictors used in the model are related to information security.
Dependent variable: all variables of Table 9 (H1 Instrument) along with variables $H_2 6, 7$ and 8.
Predictors: years of experience, years in current job position, experience of security incident, security-operations tradeoff today, closeness of security to business objectives today, closeness of security to business objectives in job environment, willingness to sacrifice security for speed of operations, job title, need for more confidentiality, integrity and availability measures in job environment, person who makes security decisions at work, salary, power to make independent security decisions at work.
Sample: professionals.

Specification 4 is different from the first three specifications. In this case, we considered WTP as fixed preference and we explored the influence of the expressed 'worry' of the subjects on WTP.
Dependent variable: worry about security incidents at work and worry about new unidentified security threats.
Predictors: age, gender, education, marital status, number of dependents in family, language.
Sample: professionals and students.

D Definitions

H_xy:	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.
$H_2$1 to $H_2$5:	Variables that describe comparisons of pairs of L_i lotteries.
$H_2$6 to $H_2$8:	Five-outcome lotteries with large losses; 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 $\mu = -25$.
*Scenario*1:	Experiment question in which participants chose between enhancement of either security or operability.
*Scenario*2:	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.

$SWITCHPOINT\_SEC$:	Variable that denotes a switching point of enhancing security by x% and operability by 10%, after which, operability enhancement became more attractive to the subject.
$SWITCHPOINT\_OPS$:	Variable that denotes a switching point of enhancing operability by x% and security by 10%, after which, security enhancement became more attractive to the subject.
$LOSS\_AV\_SEC$:	Variable that measures the difference between SWITCHPOINT_SEC and elicited preferences of Scenario 2.
$LOSS\_AV\_OPS$:	Variable that measures the difference between SWITCHPOINT_OPS and elicited preferences of Scenario 2.
RiskAversionHx_y:	Variable that measures the difference between participants' WTP and the expected value of lottery H_xy.