# Economic methods and decision making by security professionals

**Adrian Baldwin (HP Labs), Yolanta Beres (HP Labs), Geoffrey B. Duggan (University of Bath), Marco Casassa Mont (HP Labs), Hilary Johnson (University of Bath), Chris Middup (Open University), Simon Shiu[1] (HP Labs)**

## Abstract

*Increasing reliance on IT and the worsening threat environment mean that organisations are under pressure to invest more in information security. A challenge is that the choices are hard: money is tight, objectives are not clear, and there are many relevant experts and stakeholders. A significant proportion of the research in security economics is about helping people and organisations make better security investment and policy decisions.*

*This paper looks at the impact of methods based on security economics on a set of decision makers. Importantly, the study focused upon experienced security professionals using a realistic security problem relating to client infrastructure. Results indicated that the methods changed the decision processes for these experienced security professionals. Specifically, a broader range of factors were accounted for and included as justifications for the decisions selected. The security professional is an (important and influential) stakeholder in the organization decision making process, and arguably a more complete understanding of the problem is more suitable for persuading a broader business audience.*

*More generally the study complements all research in security economics that is aimed at improving decision making, and suggests ways to proceed and test for the impact of new methods on the actual decision makers.*

## 1. Introduction

The growing threat environment and increasing reliance on IT mean that security investment and policy decisions are becoming more difficult and more business critical. In large organisations, security decisions involve many stakeholders, including IT, finance, compliance, business and risk managers. This makes the decision process more complex as different stakeholders all have different knowledge, expertise, and incentives relating to security. For example, the security team normally has the subject matter expertise, but lack the business context to properly make the business case for an investment.

A key element with security decisions is the complexity of the problems. Typically, a decision to implement one or another security procedure requires the consideration of a huge range of inter-dependent factors, some of which vary in complex ways. Moreover, it is difficult to know or predict the actual impact of different choices on these factors. Extensive background knowledge about security and the company, prior experience of making similar decisions, and established standards such as ISO27000, see [ISO], help security professionals to cope with some of this complexity. Nonetheless, there is plenty of evidence indicating that even experts find it difficult to accurately trade-off multiple variables simultaneously [PB93]. This is essentially a problem of limited cognitive processing capacity – the

[1] Corresponding author – simon.shiu@hp.com

decision maker is unable to hold all of the required information in mind whilst carrying out the necessary computations.

The most common tool for supporting decision making in this way is simply a pen and paper. Writing information down lessens the amount of information being held in working memory and frees up cognitive resources. In a sense, users "download" cognition to the environment. Payne, Howes and Reader [PH01] show that users adaptively allocate resources between internal and external cognition to maximize performance.

In this paper we describe a study with twelve experienced security professionals to examine how security decisions are made and justified. Including preparation of scripts and tools, practice runs, iterations, finding appropriately experienced security professionals, and conducting the actual interviews the study took over six months to complete. The study focused on the economic utility based approach developed and described in our earlier work [BP10] together with a system modeling and simulation based on the Gnosis toolset [BG08, CB10].

Our economic utility based method aims to help decision makers identify and prioritise the trade-offs between the business outcomes of a security decision, and as a result extracts a form of utility relevant for a decision maker and/or their organisation. We start from the assumption that at least three outcomes, such as cost, productivity and security risk, trade-off against one another. The decision maker is guided through multiple steps where he/she has to prioritise the outcomes, select appropriate measures that can be used as proxies for the outcomes, and finally express the targets for these measures and the preferences in meeting them. Results from Gnosis based system modeling and simulation are then used to help the stakeholders gain a better understanding of their assumptions and to show the predicted effect that a security decision has on the selected measures and business outcomes.

The study was designed so as to examine the difference (if any) these techniques make to the security decision making process. Specifically, if and how they effect:

- the conclusions or decisions made,
- the thought process followed,
- the justifications given, and
- the confidence the stakeholder has in the final conclusions or decisions made.

The focus was upon the way our methodology and related software tools influence the security professionals as a precursor to understanding how in turn this may influence organisational decision processes. To this end, the security decision problem for the study and the possible alternative solutions were chosen to require participants to make different trade-offs between security, productivity and cost. There was not an expectation that the use of the methodology and tools should lead to any particular decision outcome to be favored. This reflects the multi-factorial and often ill-specified decision making typically undertaken by the security professionals.

This paper describes the process followed through the study, the economic based methods used, and the analysis of the results. It is organised as follows. Section 2 discusses related work; section 3 describes the economic framing and system modeling approaches used; section 4 outlines the study goals; section 5 describes the structure of the controlled study and the phases carried out by all participants; section 6 describes the economic and modeling interventions that half the participants followed; section 7 describes the data analysis and results; section 8 provides a discussion and interpretation of the results; and section 9 summarizes and draws conclusions.

## 2. Related Work

There have been many examples of economic methods being applied to improve or explain security outcomes [And01, AM06]. Of most relevance to this study is the work that proposes techniques or suggestions for how organisations should make or justify their security investments. Gordon and Loeb [GL06] describe methods for how to provide a return on investment justification for security decisions. Schneier [Sch08] further discusses the challenges of security decisions and suggests cost benefit analysis as providing a more appropriate framework for justifying business decisions. The UK Government funded collaborative project on 'Trust Economics' [TE11] provides a series of examples where economic, modeling and empirical methods are combined to improve security decision making. These include studies of USB stick policy [BC09], tool support [PM10], human factors [BS08], patching policy [IP09], de-perimeterisation [BP10] and identity management [CB10]. Related to the trust economics project and specific to the case studies and tools used in the study methodology are the examples on vulnerability and threat management [BG08] and identity and access management [CB10]. By providing an analysis of how the security professionals and decision makers actually make a decision, and how they may be influenced, this work is complementary to all of the above.

There is a large body of research on decision making and support [FM09]. Keeney and Raiffa [KR76] provide a comprehensive description of approaches to multi-stakeholder, multi-objective, multi-attribute decisions many of which are similar to the economic framing process used in our study. Security is an especially difficult and rich area for decision support. Most major security decisions involve multiple stakeholders with multiple objectives. In addition the stakeholders do not have shared understanding of the context, lack common language; have to make decisions with very little empirical data, continually changing threats, technology and best practices; and any predictions can only be made with high degrees of uncertainty.

With the study described in this paper we aim to better understand how the multi-objective decision support approaches that have been applied in other areas of decision support can be used in the security domain.

The factors in the security decision process detailed above mean that it is impossible for a security expert to precisely weight all the relevant variables and derive the optimal solution accordingly. The literature on decision making under uncertainty indicates that the way individuals cope with such problems is by employing heuristics or rules of thumb [Ka03]. These shortcuts are relatively easy to

implement and have been shown to be remarkably effective [GG96]. Nonetheless, they can also lead to systematic biases in performance.

A bias that has been well documented within the psychological literature is the tendency for people to seek information that confirms their viewpoint rather than information that contradicts their viewpoint. This bias does not refer to a conscious, deliberate attempt to selectively gather information to support a particular perspective as, for example, is undertaken by lawyers in a court case. Rather it refers to a less explicit, unwitting selection of evidence without intent to bias the conclusion (see [Ni98]).

This confirmation bias can affect decision making carried out by security professionals. Security professionals are by definition experts in security and, even where this is not the case, they are motivated to be perceived as knowledgeable about security. This expertise can lead to a high level of confidence in any initial decision made making it less necessary to pursue alternative solutions. Another consequence is that any disconfirming evidence could be challenging to their self-concept as an expert.

This confirmation bias was demonstrated using simple reasoning tasks which found that individuals tested a hypothesis by seeking information that was consistent with the hypothesis rather than seeking inconsistent information that would disprove the hypothesis [Wa66]. Studies have since demonstrated that the confirmation bias applies both when a decision has been made and prior to the decision [SF00]. Further, there is evidence that a preexisting preference can lead to the distortion of new information in favour of the preferred alternative [RM96].

Outside of the laboratory the confirmation bias has been observed within contexts as diverse as policy making during the Vietnam war [Tu84], medical diagnosis by physicians [ES78] and the development of scientific theory [Ni78]. These real world case studies illustrate that expertise within a particular area does not preclude the operation of the bias. Moreover, they show the generalizability of the bias to complex environments where decision makers must balance multiple variables against one another.

In this work we examine the impact of the confirmation bias upon security decision making and discuss how our methods aim to address this bias.

## 3. Economic Framing and System Modeling

Figure 3.1 shows the overall approach from security problem to iteration between eliciting preferences and utility through economic framing, and exploring and predicting consequences through system modeling.

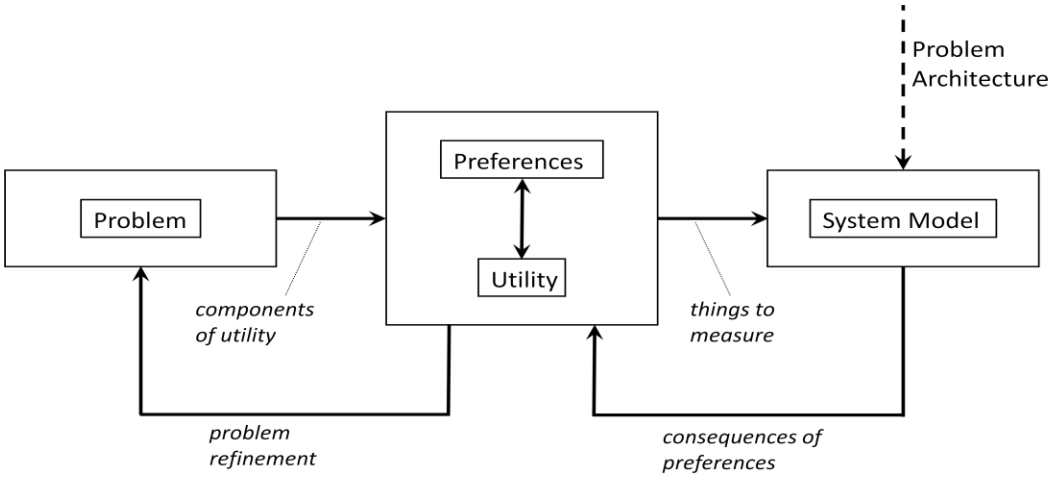


**Figure 3.1 Iterating between Economic Framing and System Modeling**

Organisations need to determine an appropriate policy, process, and technological response to the threat faced by the organisation in the context of operational requirements and security budget. Among the many attributes that must be considered are information confidentiality and integrity, system availability, assurance, and business performance. Moreover, this multi-objective, multi-attribute decision problem must be solved in a highly variable, highly dynamic environment.

There are many approaches based on security economics that address this problem. In this study we have focused on the combination of economic framing and system modeling that has been developed from a series of previous case studies including vulnerability and threat management [BG08], USB stick policy [BC09], de-perimeterisation [BP10], and identity and access management [BCS09].

In this approach economic framing is used to identify and prioritise between the multiple objectives. The framing is provided by a form of multi-criteria utility function. In practice, it is very difficult to extract a formal utility function direct from the organisation. We approach this with a multi step process that guides participants to select outcomes and preferences relevant to their organisation. This includes choosing how these outcomes can be measured, which are the most important, at which points performance in a particular outcome becomes 'intolerable' and predicting what effect the security decision could have on the multiple outcomes. A fuller exposition of the method is given in [BP10], and further discussion and examples of applying and using this style of utility function are given in [IP09,BC09].

System modeling and simulation are used to help the decision maker explore the effect the different security controls have on the selected set of outcomes. For example, following empirical work to construct a model we simulate the effect that restricting admin rights might have on reducing risk exposure to malware-type threats and model its impact on user satisfaction.

In this study system models are built and explored using the Gnosis language and tools [CM10]. Gnosis is a discrete process simulation language, and can be used to explore complex concurrent interactions of processes and resources (at locations) under numerous assumptions and stochastic conditions. This

approach is particularly useful for exploring the combined effect of various operational processes and technological solutions. Much work has been done on the mathematical foundations behind the Gnosis language and its suitability to modeling security problems [CM09].

## 4. Study Goals

The goal of this study was to explore how security economics can change the way security professionals think about and justify security decisions. In order to address this question it was necessary to conduct in-depth studies not just of the decision itself but also the process of making that decision. To maximize the validity we chose to work with experienced security decision makers and to examine the process by which a decision to a presented problem is made. This strategy meant that our analyses were primarily qualitative and considered each participant's individual decision process for similar approaches see [EL96, LK01, Pa76].

Overall twelve current security professionals were involved in the study. Careful selection ensured each had many years of experience of security decision making in a mixture of consulting, operational and policy roles. Most had significant experience of security in public sector and financial services, with others having experience of commercial, retail telecoms and energy organisations.

It is unusual and therefore novel to have this amount of time with a significant number of professionals. The intention was to explore if they could use the information provided, whether it was deemed relevant and whether it figured in their justifications. For example, do the techniques lead to professionals asking more questions, making more causal links, and being aware of different trade-offs. If so then this small-scale study will have explored and illustrated some of the value of our economic approaches.

Each participant was presented with the same security problem and four possible security decision options. The decision problem involved making improvements in the vulnerability and threat management for client infrastructure in an organisation. The organisation was fictitious with a number of key attributes established. These included company size, regulations that applied, structure and size of the IT workforce, history of security incidents, IT architecture, critical business processes, mobility of workforce and so on.

Four decision options were presented: (1) investing to improve software patching, (2) buying and deploying client based host intrusion prevention technology (HIPS), (3) introduce changes in when users have admin privileges, and (4) do nothing. Our previous work in the area of vulnerability and threat management with several large organisations [BG08] provided confidence that the problem selected was realistic and representative, and meant we can call on our experience for how the security stakeholder's opinion and role affects the organisation decision process.

Even though the decision problem is restricted to four investment options, it is still a complex multi-objective problem with a high degree of uncertainty with how well the options might meet the organisations business objectives. From one point, each security investment address different types of security risk. For example, investing in patching might ensure the operations team does a better job of

getting infrastructure sufficiently patched within policy deadlines; investing in HIPS might provide quicker protection for a significant class of threats; and investing in lockdown may prevent privilege escalations that in turn reduce the spread of malware. Each option is also likely to have a different impact on productivity, costs, compliance. For example, doing nothing presumably does not reduce security risk, but does have the advantage of not affecting capital or operational costs, and of not impacting productivity of users, IT operations or business processes.

Though all participants were introduced to the problem in the same way, only half of the participants (participants in each group were matched in terms of experience) were guided to use the economic utility based methodology (with the help of specifically designed software tools) and the results from simulations of the system model. This group (the intervention group) was challenged to think through preferences for how the different strategic outcomes trade-off, and to use a system model and associated results to help explore the effects of the four decision options on these outcomes. The other half acted as a comparative control group, and was asked to make the decisions without the help of the extra tools. This second group functioned as a control or baseline against which to compare the intervention group's performance.

For reasons of brevity in this paper, types of questions and justifications for each participant were simply aggregated to highlight differences between the groups. Further analysis is being done to look at the trajectory of each individual participant question/justifications, to see the different processes within individuals and to discern any patterns within and between groups, and for the intervention group whether and where information from the interventions were mentioned or used in later questions and justifications.

An interview lasted 30-45 minutes for a participant from the control group, and 60-75 minutes for participants from the intervention group. Each interview had two subject matter experts; one running the interview and ensuring consistency of protocols for introducing information, tools and moving through the various stages, and a second acting as a member of a security team of the fictitious organisation as well as an expert of the tools used (when required), answering all questions that the participants asked. All the sessions were recorded and transcribed, and the tools captured choices made and screens viewed.

## 5. Study Structure and Common Phases

Figure 5.1 shows the scripted phases followed by the intervened and control groups. The problem description, decision options, choice and justification and introspection phases were the same for each group. The only difference for the intervened group was that we included phases 5a and 5b where the participants worked through economic framing and system modeling. Below we describe details relevant to each of the main phases shown in the diagram.

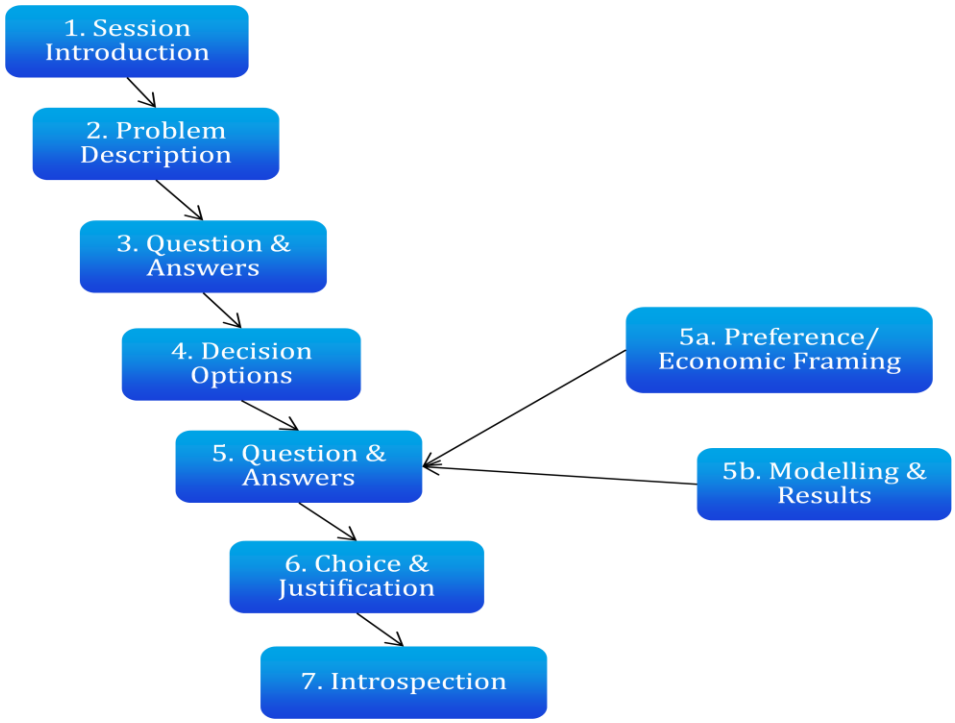


**Figure 5.1. Schematic of study phases for each group**

**Study Phase 1. Session Introduction**

All participants were given a standard introduction to the session. This included an introduction to the experimenters, an outline of the role they were expected to play, and the incentive for doing this to the best of their ability. In addition they used a web based tool to record their security experience.

**Study Phase 2. Problem Description**

All participants were given a standard written description of the problem. Essentially they were presented with the scenario that they had just joined the security team of a company, that the CISO was concerned about protections on client infrastructure, and was looking for a recommendation for whether and how to proceed.

**Study Phase 3. Question and Answers**

At this point they were encouraged to ask any questions they thought relevant to the problem and the expert used a script sheet to provide consistent answers. For example, there were scripted answers for the type of business, compliance culture, history of incidents, the different needs, behaviours and expectation of staff in different functions, and so on. If a subject asked a question that was not part of the script, the expert made up an answer and this was recorded to ensure consistency in future interviews.

All participants were given 10 minutes to ask questions at which point if they had not asked or been told about the patching policy and processes they were given a sheet describing this. This allowed us to explore what order participants sought information, but ensured all participants had certain key information that we deemed relevant to the problem, before being asked to provide a decision.

**Study Phase 4. Decision Options**

To simplify the problem and the results analysis the participants were given 4 discrete decision options:

1. process and technology investment to improve patching of client devices,
2. investment to deploy host based intrusion prevention (HIPS) technology on all clients,
3. rolling out a policy to lockdown all clients so that only the IT department would have administrative privileges, or
4. to do nothing.

As mentioned in Section 3, the problem is based on previous real world case studies so these options are known to be realistic and representative [BG08], [BP10]. Key to this study the implications of these options are well known to security experts and they represent interesting and distinct trade-offs, concerns and issues, as discussed in more detail in [BG08].

All participants were invited to provide an initial view at this point.

**Study Phase 5. Intervention**

This phase was only followed by the intervention group. The two parts (5a and 5b) and described in section .

**Study Phase 6. Choice and Justification**

Whilst we welcomed discussion on how combinations of other solutions, and how combination or sequenced strategies would make sense, we were clear we wanted a preferred option. Once they had given their preferred option all the participants were asked to fill out a form for each option describing pro's and con's, and a confidence level (1-7 Likert scale).

**Study Phase 7. Introspection**

After the participants had completed their justifications they were encouraged to introspect on how they had solved the problem, why they sought certain information, and how it had been used. For the intervened group we asked what, if any difference they felt the tools had made to their thinking.

## 6. The Intervention Phases

This section describes the way the intervention group was exposed to the economic framing and system modeling approach described in section 3. These interventions are specifically designed to deal with the challenges decision makers have dealing with multiple factors and predicting the actual impact of

choices. In both phases a number of tools and materials were used and so these are explained ahead of the actual phase description.

## 6.1 Tool support for Economic Framing

This phase involved a complex workflow where participant's choices about business priorities would affect which proxy measures and trade-off graphs would make sense to display and process. A preference elicitation tool was used to capture, process and display information in support of this process. The high level overview of this tool is provided in appendix 1. Figure 6.1 provides various screenshots of this tool, for the different steps, including: the initial questionnaire to gather information about the interviewed person; elicitation of outcomes of relevance, elicitation of relevant proxy measures; and trade-off graphs between chosen proxy measures.

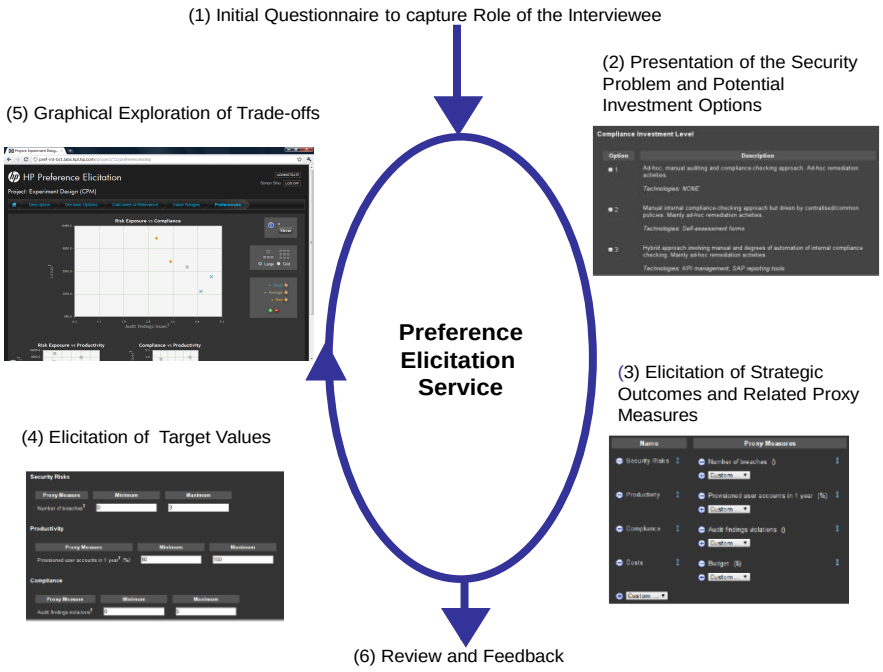


**Figure 6.1 - Screenshot from economic framing tool**

**Phase 5a. Economic Framing (see figure 5.1)**

At the start of the economic framing phase participants were shown the strategic outcomes screen. This suggested significant business outcomes such as compliance, (security) risk, (business or user) productivity and cost. They were encouraged to think of other components we may have missed, and then to reduce to 3 main components that seemed most relevant to the business decision they were making.

Participants were then shown a series of metrics, or proxies for the chosen components. For example:

- would the cost component best be measured by impact on operational cost, capital cost, some combination, or something else?
- Would the productivity component best be measured by user satisfaction, infrastructure downtime, or something else?

They were asked to select the one they thought would serve as the best proxy for measuring performance of the component.

The tool took these inputs and generated a series of trade-off graphs, see figure 6.2. In the figure the components are risk and compliance, with proxy measure being financial losses and annual audit findings. Clearly low losses and low numbers of audit findings is desirable (bottom left of the graph) and high losses, high number of findings is undesirable (top right of the graph). Once the participants understood this they were asked discuss (and mark) how they felt relatively about the areas that trade off, i.e. is it better to be in the lower right (low losses, but higher number of audit findings), or upper left (higher losses, but low audit findings).

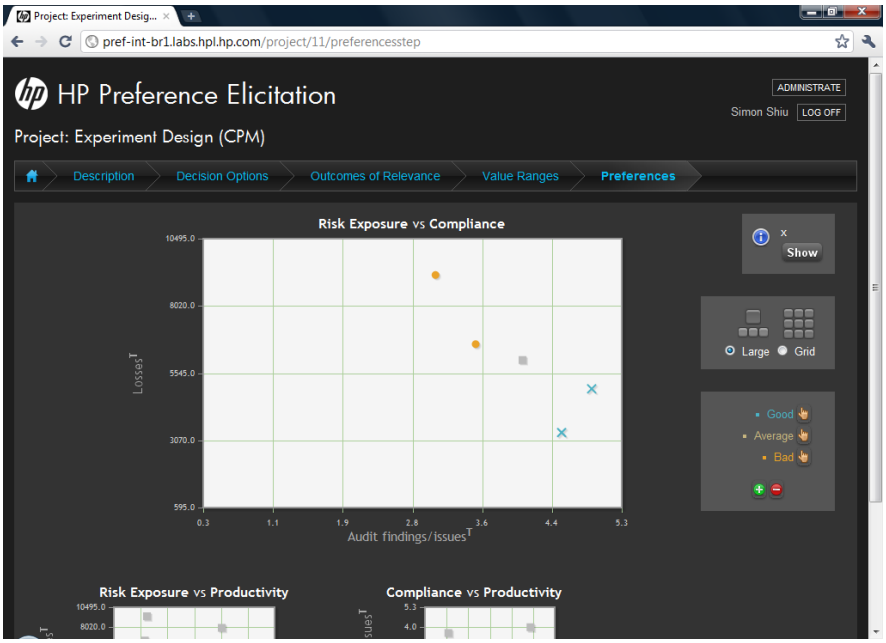


**6.2 Screenshot of trade-off graphs from the preference elicitation tool**

This ensured that we could record the prioritized preferences expressed, and have some confidence they properly considered both strategic outcomes and the way they trade off.

## 6.2 Tools and Material Support for System modeling and Results Analysis

In previous case studies the stakeholders were involved in the empirical work to construct a model. This was not possible for this study, instead we used a model developed to support a similar problem. Figure 6.3 shows the system diagram of the model used in the study. The boxes represent processes, event generators and decision points, and show the paths that can be taken once a vulnerability is discovered.

A simulation allows these paths to run concurrently, and to collect metrics such as the time taken from vulnerability discovery, to “some” mitigation being in place.

*The red section* within the model represents the external environment of vulnerabilities, exploits, malware and patches. This includes statistics observed from the external world. *The blue section* represents the security controls and processes within the organisation, such as patching and AV Mitigations. Patching is split into testing and deployment sub-processes and further separated to capture accelerated patching when the vulnerability is deemed critical (decisions for when to accelerate are represented by the complex set of logic gates) and an emergency patching when it is critical and no other protections are in place. The *additional red boxes* represent the additional security investments in HIPS and lockdown that are added to the model for some simulations. Investment in patching is represented by making different assumptions about how long patch test, deploy and acceleration would occur. A more detailed description of this model can be found in [BG08].

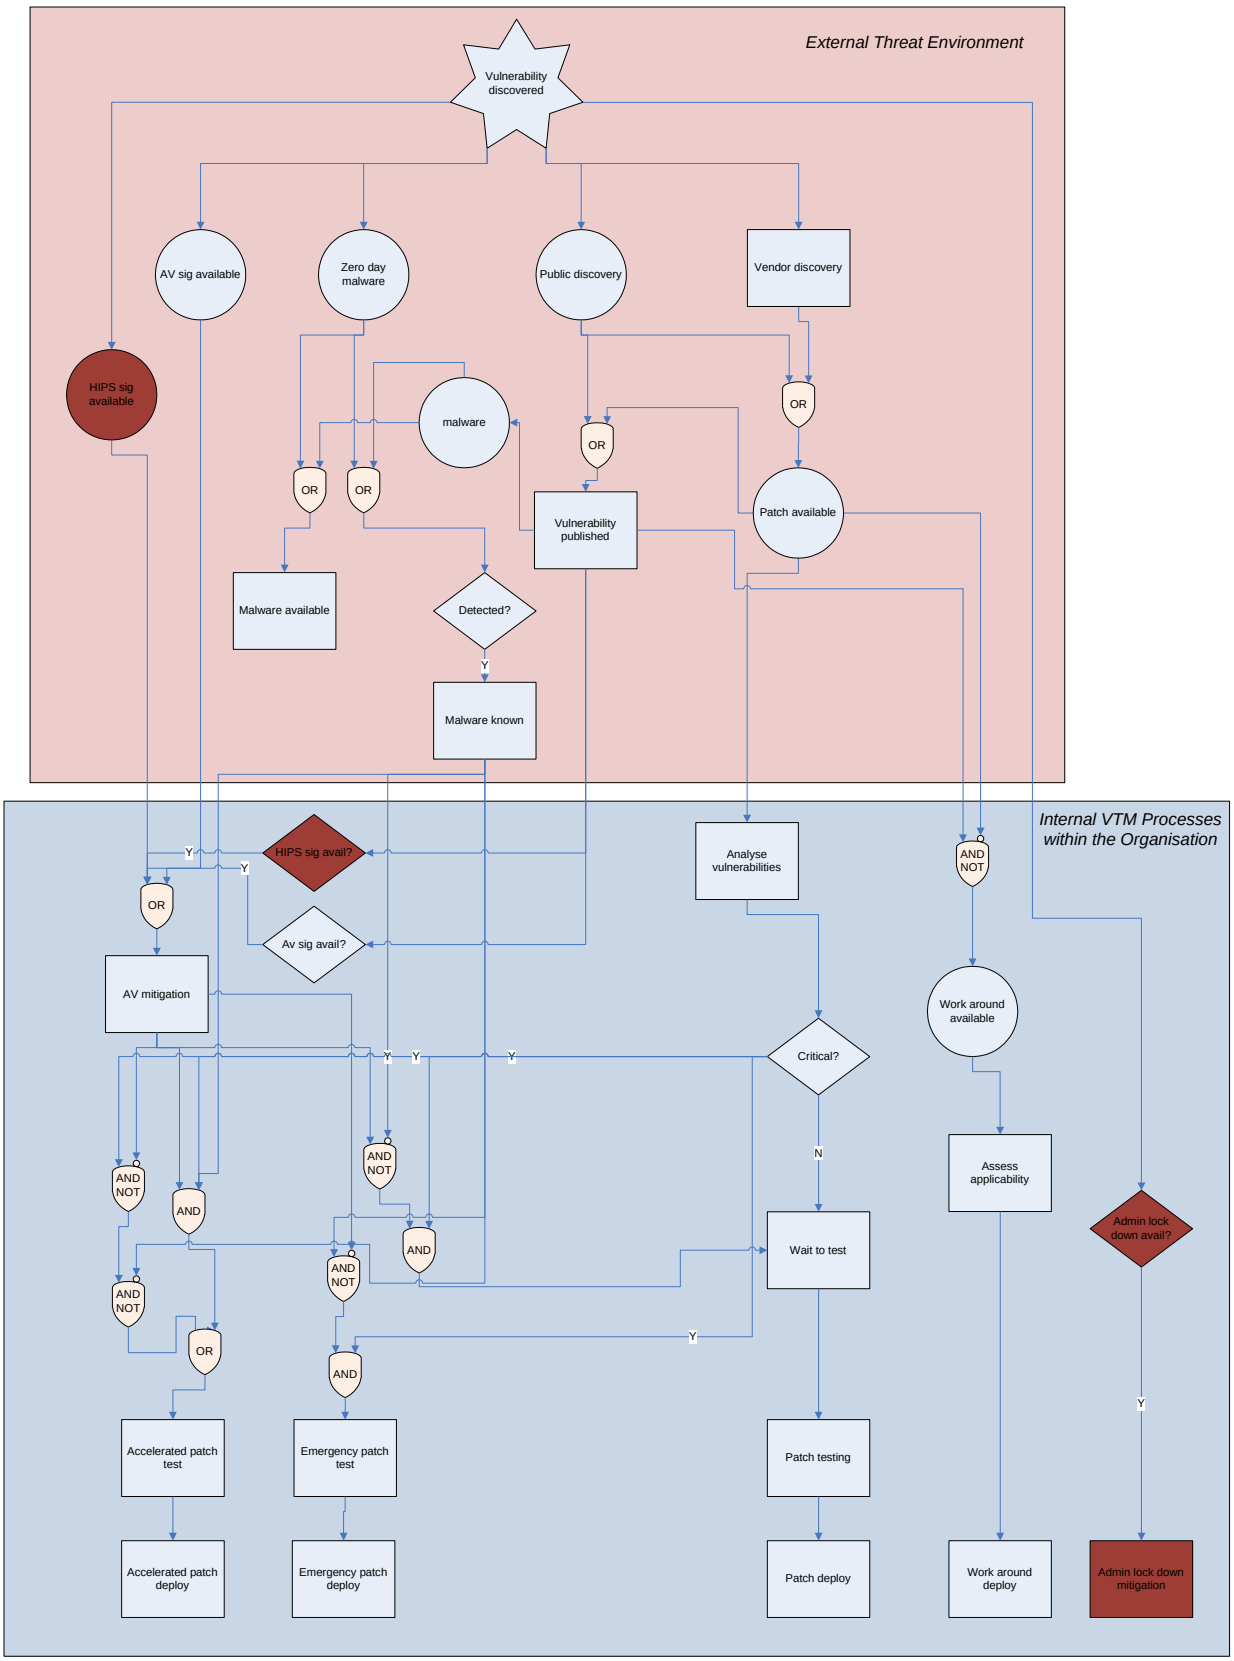


**Figure 6.3 - Schematic of the system model used for the study**

The model was configured to produce statistics relating to the proxies discussed in the economic framing. Based on previous experience we added details about the effect on user satisfaction, impact on the productivity of operations staff, and associated costs. Overall we had data for nine measurements that could be contrasted for each of the four decision options:

- *Machine days exposed to know malware* - # of machine and # days each of them is exposed from the time that malware is known in the wild
- *Exposure window* – day's workstation environment is exposed (not mitigated) from vulnerability discovery time
- *Malware infection* - # of machine infections
- *Helpdesk calls* – number of helpdesk calls per each vulnerability.
- *Productivity of operational staff* – number of hours spent by operational staff doing security tasks
- *User satisfaction* – level of user satisfaction between 0-4.
- *Policy violations* - % of vulnerability cases where mitigations took longer to be deployed than policy dictated timeline
- *Capital cost* – one off dollar value
- *Operational cost* – dollars spent yearly.

We created a results tool to mine this data and produce comparative views. For example figure 6.4 shows a screenshot from this tool which compares results for cost, exposure and staff productivity for 3 of the decision options.

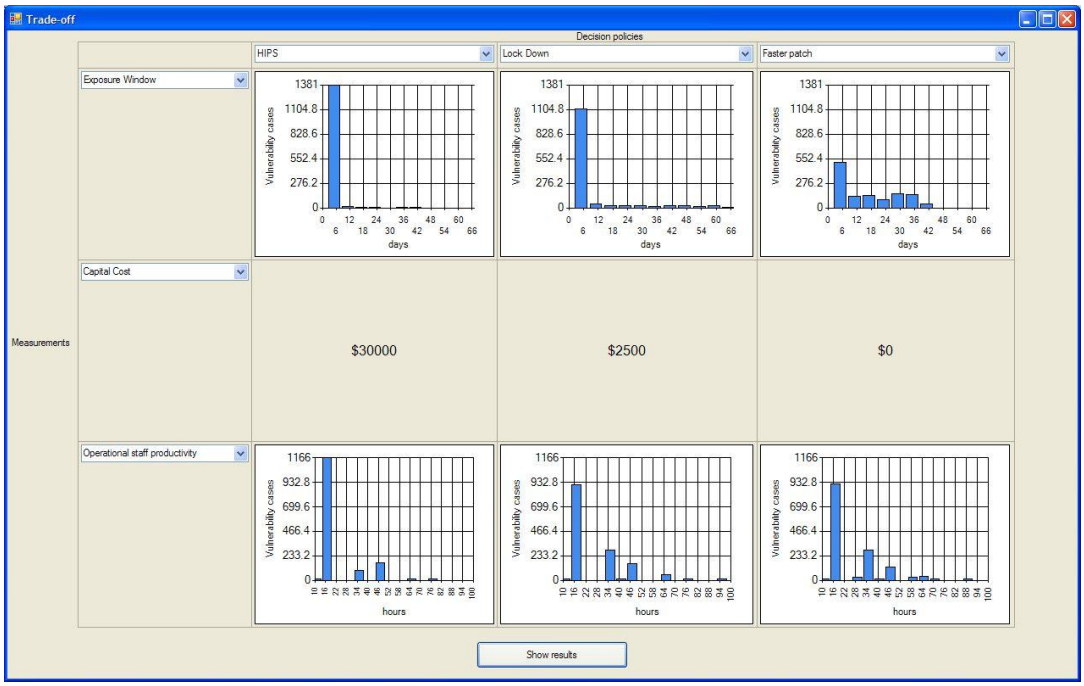


**Figure 6.4 Screenshot of the system model results tool**

**Study Phase 5b. Modeling and Results (see figure 5.2)**

After economic framing the participants were shown a schematic of the system model, see figure 6.3. The interviewer explained that domain experts had verified that the model was a reasonable representation of the system, that simulations could be run to see the effect of the different decision options under a range of conditions and that the results of simulations provided reasonable evidence of expected behavior and outcomes in each of the simulated situations.

The participants were introduced to the results tool and shown results for options and proxies they had prioritized earlier in the interview, see figure 6.4. They were then able to select different proxy and decision option combinations so they could compare multiple outcomes from different choices.

# 7. Results and Data Analysis

This was a small-scale study conducted on twelve security professionals, which means the results presented here cannot and should not be taken as statistically significant. The results point to differences that might be replicated in studies with larger numbers. This section shows and describes some differences that were observed, and the following section provides theoretical explanations for these differences.

## 7.1 Questions and Justifications

Some statistics and examples of the data collected are provided in appendices 2 and 3. To ensure an objective approach to the analysis, all the questions and justifications were transcribed and presented in a random order. Two security experts and a cognitive scientist then independently categorized each of the questions and justifications. Any disagreements were resolved through discussion. The categories were:

Cost - meaning the question/justification was about finance or budget considerations
Compliance – meaning the question/justification was about regulation constraints
Productivity – meaning the question/justification was about effect on business or user productivity
Evidence – meaning the question/justification was about historical data, events, or incidents
Security – meaning the question/justification was about improving security

The remaining questions and justifications did not fit into the above, and did not form any other grouping, and so were all labeled “other”.

The proportion of questions and justifications in each of the categories are given in figures 7.1 and 7.2. For the questions in Figure 7.1, the distribution across the categories was similar in both groups. This is unsurprising given the questions were asked prior to the intervention. The majority (just over 63% in both groups) of questions referred to security issues and there were very few questions about costs or productivity implications.

By contrast there was a clear difference between the groups in the proportion of justifications allocated to each of the categories. Figure 7.2 shows that security remained dominant in the control group (with 51% of justifications), this dropped considerably (to 36% of justifications) in the intervention group.

Moreover the emphasis on cost and productivity was considerably higher in the intervention group than the control group. These results suggest that the economic framing and system modeling led to a broader range of reasons being produced to justify the decision in the intervention group.

To ensure that this greater breadth in justifications was not at the expense of depth in reasoning the justifications were scored according to depth of explanation. For example, simply mentioning cost would score 1, explaining the impact in terms of operational costs would score 2, whereas explaining the cost impact in comparison with the gain/loss of productivity would score 3. As with the other categorization this was carried out by security experts. Based on this analysis more justifications were provided by the intervention group (81) than for the control group (68), and there was a small increase in depth of explanation per justification, intervention Group: M = 1.91, SD = 0.24; control Group: M = 1.8, SD = 0.34.

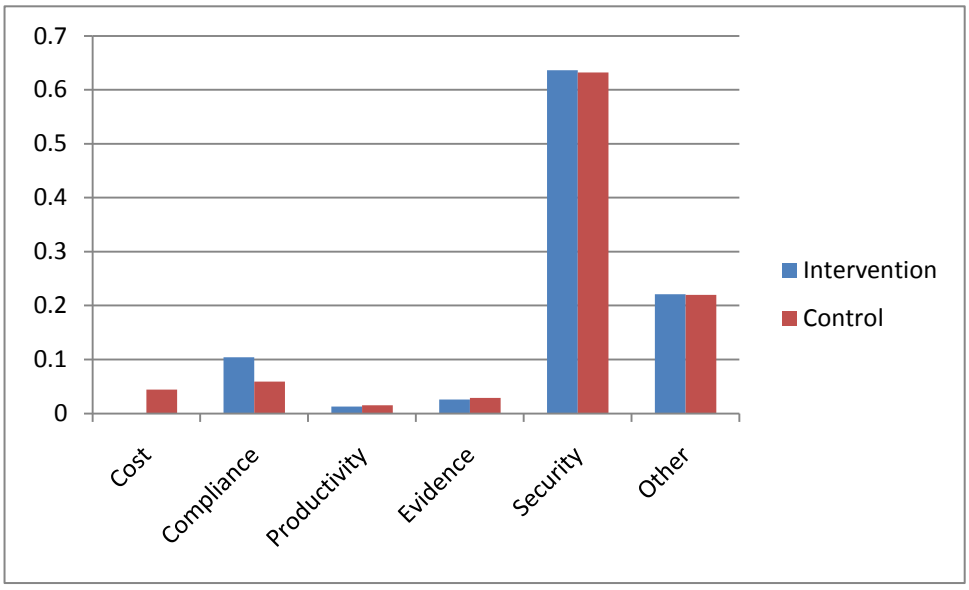


**Figure 7.1 – Proportion of questions in each category for intervention and control groups**

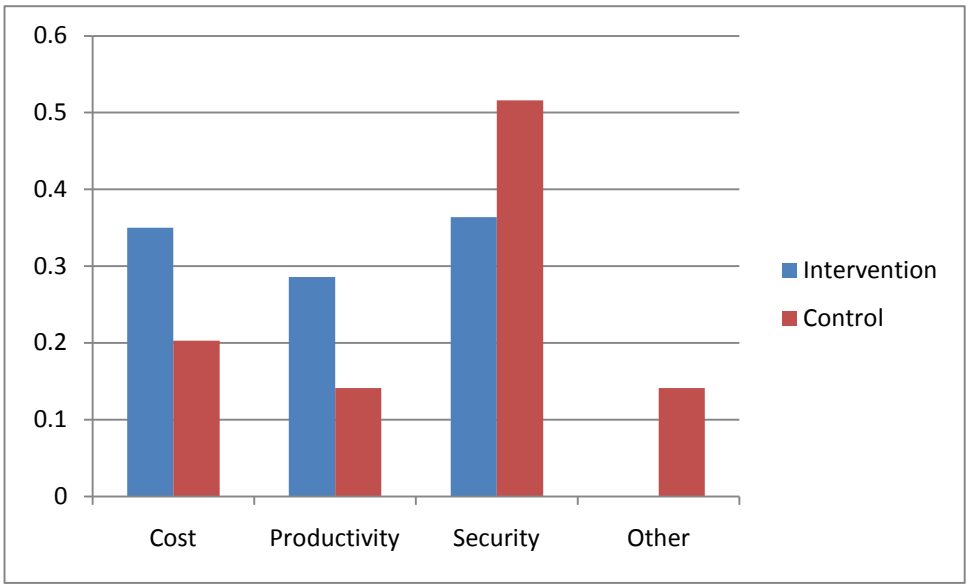


**Figure 7.2 - Proportion of justifications in each category for intervention and control groups.**

## 7.2 Results of economic framing

The participants in the intervention group were asked to drop a factor, and then express on a graph how they felt about relative trade-offs between the others. This allowed us to derive an ordered set of priorities for each participant from highest priority (1) to lowest priority (4). This found that Security risk had the highest priority (M = 1.50, SD = .77), followed by Productivity (M = 2.17, SD = .82), Cost (M = 2.75, SD .76) and then Compliance (M = 3.58, SD = 1.02). The preferences show that although Security was considered the most important factor the economic framing encouraged all participants to also consider both cost and productivity.

When the justifications were categorized into the utility components (i.e. each pro and con was labeled as best belonging to either cost, risk, compliance or productivity) there were many more cost and productivity reasons provided by the intervention group. In most cases the justifications reasonably represented the priorities expressed, although in one case the opposite occurred (i.e. where compliance and risk were expressed as the priorities the justifications were actually expressed only in terms of cost).

## 7.3 Results of choices and Likert scale confidences

Decision outcomes were similar in both groups with 3 participants selecting HIPS and 3 selecting lockdown in the intervention group. In the control group 3 selected lockdown, 2 selected HIPS and 1 selected patching. Mean confidence for each of the 4 options in both groups is given in Figure 7.3. There was slightly higher confidence for both the HIPS and Lockdown options in the intervention Group. Mean confidence for the selected option was 5.83 (SD = .75) in the intervention Group and 5.50 (*SD* = 1.05) in the control group.

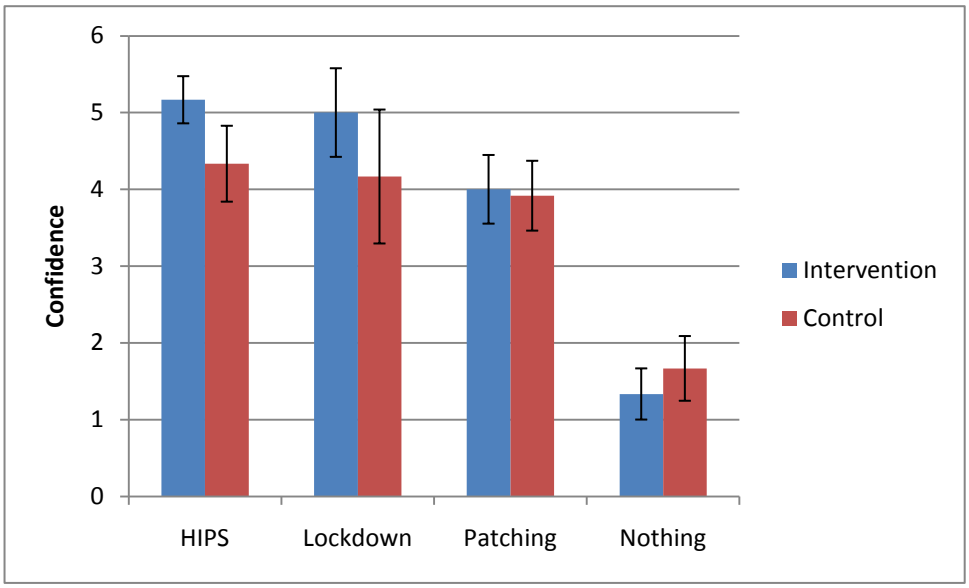


**Figure 7.3 - Mean confidence and standard error for each decision option.**

### 7.4 Introspection results

The main themes from the introspection were that intervened participants valued the results tool but not the economic framing. Specifically, for the results tool they liked the quantitative results of simulations and the ability to easily compare results of different options. Conversely for the economic framing, they typically reasoned that they were already well aware of the main outcomes and trade-offs. This makes the result that the intervention did seem to affect the richness and range of their justifications more interesting.

## 8. Interpretation of the Data Analysis and Results

The results from the study provide a clear indication that the decision processes of security professionals were affected by the economic framing and system modeling tools. Initially, participants gathered information about the problem by asking questions and in both groups, the majority of questions focused upon security related factors. Having made a decision, participants were required to explain their reasoning for each of the decision options. In the control group the majority of these justifications related to security issues. However, in the intervention group, the type of justifications produced for the decision focused upon factors such as cost and company productivity as well as security related factors.

*Confirmation Bias*

The high proportion of security related questions asked in both groups suggests that security professionals typically focus on factors that are related to their expertise. This confirms our initial view that the confirmation bias described in section 2 affects the decision making carried out by security professionals.

We can also observe that the economic utility based approach encouraged the alternative factors to be considered which worked against this confirmation bias. Firstly, it shifted the focus away from purely security factors and secondly, by representing different trade-offs explicitly to security professionals and by encouraging them to consider interactions between variables, there was greater opportunity for disconfirming evidence to be evaluated.

*Cognitive Dissonance*

Interestingly, despite the differences between the two groups in the reasons cited to support their decisions, the intervention group participants claimed that their decision making was unaffected by the economic framing and system modeling. At one level this is unsurprising, it is well documented that self-reports can be unreliable measures of behaviour – people have limited access to their internal cognitive processes [NW77]. Indeed the reason that process measures were included in conjunction with self-report measures was to compensate for potential inaccuracies in reporting. Nonetheless, to understand the effectiveness or otherwise of our tools it is helpful to understand why this discrepancy exists.

A theoretical account for the confirmation bias can also provide an explanation for the failure to acknowledge an effect of the tool despite the shift in reasons provided for the decision. Cognitive dissonance [Fe57] refers to the discomfort experienced when two conflicting cognitive elements (e.g. beliefs) are held by the same person. One way to reduce this discomfort is to alter one of the beliefs to make it consistent with the other belief. Importantly, it may sometimes be easier to modify an accurate belief than an inaccurate one. This can be seen with post-hoc rationalizations when an expected outcome is not attained, for example, “I didn’t try my hardest which was why I didn’t win the race”. Some beliefs are more integral to an individual’s self concept or worldview than others and they will be more resistant to change. Thus, to reduce dissonance an alternative cognitive element is altered.

In our study, the participants were experts and were employed as such. Further, they were required to make a decision prior to using any tools. Thus, any additional insights provided by using the tools could be interpreted by the participants as a challenge to their internal perceptions of expertise. One way to reduce this cognitive dissonance is to minimize the perceived worth of the tools. We do not mean to imply security professionals will not adjust their decisions or are not open-minded. Indeed, we would argue that the shift in justifications for their decision implies that our participants did account for the tools within their reasoning. Rather, we wish to emphasize that they might not be fully aware of any benefits gained from using the tools.

Notwithstanding this, it is possible to conjecture that cognitive dissonance could negatively affect the quality of decision making by security professionals. Because the correct solution to a security problem is often open to debate it is difficult for a security professional to receive accurate feedback on decisions. Where there is not clear evidence against a decision the desire to reduce cognitive dissonance can mean counterevidence is understated. (See [GP97] for evidence of similar behaviour by mutual fund investors.)

The additional tools used by the intervention group can be seen as a solution to this as they provide information about alternative decision options. The tools lessen any confirmation bias the security professionals might have by providing both confirming and disconfirming information.

*Effect of results from system modeling*

The other noteworthy finding from the introspection was that the system modeling tool was valued as a decision aid. This supports our contention that providing external information could support the internal processing carried out by security professionals. Perhaps more importantly, it also indicates that the external support provided was in an appropriate form and was relevant for the processing required to address the problem.

*Generalizability of findings and future research*

The actual decisions made in the two groups did not differ as a result of the interventions, however, we do not view this as problematic. The problem and the alternative solutions were chosen to be representative of actual security decisions and the focus was on providing sufficient complexity to enable an understanding of the process both with and without economic interventions. We were primarily interested in the gathering of information and any subsequent reasoning which was why the questions asked and the justifications provided were of interest. Future work could focus on the actual decision by devising problems that specifically discriminated between different decision criteria yet were equivalent on other criteria. However, here our focus was upon capturing the complexity of a typical security decision where multiple attributes were traded off against each other.

As explained in section 4, our goal was to look at the way actual security decision makers solved realistic security problems and to investigate how the economic framing and system modelling tools affected this process. To this end, we have provided a theoretical explanation for the differences in behaviour associated with using the tools. In-depth study of relatively small participant sample sizes is often a richer and more fertile method for theoretical development than larger more evaluative approaches [Sa03]. We feel the findings and ideas suggested by our study reinforce this conclusion and our methodological approach. Nonetheless, we recognise the value of complementary studies that build upon the results here and generalize the conclusions to alternative problems and larger populations.

*Timing of the introduction of the tools*

The timing of the intervention also seems an important factor. Our participants were allowed to ask questions about the problem and then asked to make a decision prior to the intervention. This facilitated a controlled comparison across groups, however, it would also be interesting to study the decision process when participants were provided with the tool at the same time as the problem. This would enable our tools to be integrated with the information gathering phase of the decision. The introspection results suggested that many of the selections were based on knowledge of best practice. Providing our tools from the start would mean their benefits could be realized before or during the test for best practice. Of course, an alternative view is that best practices represent a reasonable way forward in most cases, and that our tools should be reserved for broader questions.

*Multiple Stakeholders*

This experiment focused on the security (expert) stakeholder. Our tools are designed to allow multiple stakeholders with different levels of expertise and accountability to work through and share their

priorities, preferences and assumptions. Different organisations will have different structures and processes in place to ensure due diligence on these kinds of decisions. In some cases, a risk committee may handle policy changes, whereas the CIO will make IT investment choices. Intuitively, we expect any improvement by the security professional in their understanding or ability to explain a decision should help in all these cases. However, these observations suggest further work is needed to investigate the impact of our tools upon non security stakeholders. A simpler and smaller study would be to explore whether the broader justifications are more convincing to the non-security stakeholders. A more ambitious challenge is to design a study to look at the effect our tools (or any security decision making method) have on the range of organisation decision processes.

## 9. Conclusions

Organisations' increasing reliance on IT, and the hostile threat environment mean that it is crucial to make good security investment and policy decisions. The area is complex, changing and has multiple stakeholders so making good decisions is likely to remain difficult. Security economics provides many useful approaches and techniques to improve this decision process. Trust economics [TE11] is an example project in this realm.

This study has looked at the impact of economic framing and system modeling on a set of decision makers. Crucially, the study focused upon experienced security professionals using a realistic security problem. Integrating findings from the decision-making literature into this security context has highlighted the potential for security professionals to favour information that confirms their initial viewpoint over information that does not. By externally representing alternative aspects of the problem and trade-offs between different factors our method can thus support decision making.

Results indicated that the interventions changed the decision processes for these experienced security professionals. Specifically, a broader range of factors were accounted for and included as justifications for the decisions selected. The security professional is one (important and influential) stakeholder in the organisation decision making process, and arguably the richer arguments are more suitable for persuading a broader business audience.

More generally the study complements all research in security economics that is aimed at improving decision making, and suggests ways to proceed and test for the impact of new methods on the actual decision makers.

## References

[AM06] R Anderson and T Moore, The economics of information security, Science314:610-613, 2006.

[And01] R Anderson, "Why information security is hard: An economic perspective", in proc 17th annual computer security applications conference (ACSAC) 2001

[BC09] A Beautement, R Coles, J Griffin, C Ioannidis, B Monahan, D Pym, A Sasse, M Wonham, "Modelling the Human and Technological Costs and Benefits of USB Memory Stick Security", in Managing Information Risk and the Economics of Security, Springer 2009

[BCS09] Adrian Baldwin, Marco Casassa Mont, Simon Shiu - Using Modelling and Simulation for Policy Decision Support in Identity Management, IEEE 10th Symposium on Policies for Distributed Systems and Networks, IEEE Policy 2009 Symposium, 20-22 July, London, 2009

[BG08] Y Beres, J Griffin, S Shiu, M Heitman, D Markle, P Ventura, "Analysing the performance of security solutions to reduce vulnerability exposure windows, in annual computer security applications conference (ACSAC), 33-42, CA IEEE 2008

[BP10] Y. Beres and D. Pym and S. Shiu, "Decision Support for Systems Security Investment", in Network Operations and Management Symposium Workshops (NOMS Wksps), 2010 IEEE/IFIP ,

[BS08] A Beautement, A Sasse, M Wonham, "The Compliance Budget: Managing Security Behaviour in Organisations." in New Security Paradigms Workshop (NSPW) 2008, Plumpjack Squaw Valley Inn, Olympic, California, USA. 22-25 September 2008

[CB10] M. Casassa Mont and Y. Beres and D. Pym and S. Shiu, "Economics of Identity and Access Management: Providing decision support for investments", in Network Operations and Management Symposium Workshops (NOMS Wksps), 2010 IEEE/IFIP

[CM09] M Collinson, B Monahan, D Pym, "A Logical and Computational Theory of Located Resources." in Journal of Logic and Computation, 2009. (in press) DOI: 10.1093/logcom/exp021.

[CM10] M Collinson, B Monahan and D Pym, "Semantics for Structured Systems Modelling and Simulation." in Proc. Simutools 2010, ACM Digital Library and EU Digital Library.

[EL96] K. A. Ericsson and A. C. Lehmann. "Expert and exceptional performance: Evidence of maximal adaptation to task constraints." Annual Review of Psychology, 47, 273-305. 1996

[ES78] A. S. Elstein, L. S. Shulman and S. A. Sprafka. "Medical problem solving: An analysis of clinical reasoning." Cambridge, MA: Harvard University Press. 1978.

[Fe97] L. Festinger. "A theory of cognitive dissonance." Stanford, CA: Stanford University Press. 1957

[FM09] S. French, J Maule, N Papamichail, "Decision behavior, analysis and support" Cambridge University Press, 2009

[GG96] G. Gigerenzer and D. Goldstein. "Reasoning the fast and frugal way: Models of bounded rationality." Psychological Review, 103, 650-669 1996.

[GL06] L.A. Gordon and M.P. Loeb, Managing Cybersecurity Resources: A Cost-Benefit Analysis. McGraw Hill 2006.

[GP97] W. M. Goetzmann and N. Peles. Cognitive dissonance and mutual fund investors. The Journal of Financial Research, 2, 145-158. 1997

[IP09] C Ioannidis, D Pym, J Williams, "Investments and Trade-offs in the Economics of Information Security." in Proc. Financial Cryptography and Data Security 2009, LNCS 5628: 148-162, Springer, 2009

[ISO] ISO 27000 series of standards for information security and security management, see http://www.27000.org/

[Ka03] D. Kahneman. "A perspective on judgment and choice: Mapping bounded rationality." American Psychologist, 58, 697-720. 2003

[KR76] Keeney, R. L. and Raiffa, H. (1976). Decisions with Multiple Objectives: Preferences and Value Tradeoffs. Wiley,New York. Reprinted, Cambridge Univ. Press, New York (1993).

[LK01] R. Lipshitz, G. Klein, J. Orasanu and E. Salas. "Taking stock of naturalistic decision making." Journal of Behavioral Decision Making, 14, 331-352. 2001

[Ni98] R. S. Nickerson. “Confirmation Bias: A ubiquitous phenomenon in many guises.” Review of General Psychology, 2, 175-220. 1998.

[NW77] R. E. Nisbett and T. D. Wilson. “Telling more than we can know: Verbal reports on mental processes.” Psychological Review, 84, 231-259. 1977[PM10] S Parkin, A van Moorsel, P Inglesant, A. Sasse, “A Stealth Approach to Usable Security: Helping IT Security Managers to Identify Workable Security Solutions.” in the Proceedings of the New Security Paradigms Workshop (NSPW) 2010, Concord, MA, USA, 2010.

[Pa76] J. W. Payne. “Task complexity and contingent processing in decision making: An information search and protocol analysis.” Organization Behavior and Human Performance, 16, 366-387. 1976

[PB93] J.W. Payne, J.R. Bettman, and E.J. Johnson. “The adaptive decision maker.” New York: Cambridge University Press. 1993

[PH01] S. J. Payne, A. Howes and W. R. Reader. “Adaptively distributing cognition: A decision-making perspective on human-computer interaction.” Behavior and Information Technology, 20, 5, 339–346. 2001

[RM96] J. E. Russo, V. H. Medvec and M. G. Meloy. “The distortion of information during decisions.” Organizational Behavior and Human Decision Processes, 66, 102-110” 1996

[Sch08] B Schneier, Security ROI, in Schneier on Security blog. 2 Sept 2008, see http://www.schneier.com/blog/archives/2008/09/security_roi_1.html

[Sa03] P. M. Salkovskis. “Empirically grounded clinical interventions: Cognitive-behavioural therapy progresses through a multi-dimensional approach to clinical science. Behavioural and Cognitive Psychotherapy, 30, 3-9. 2003.

[SF00] S. Schulz-Hardt, D. Frey, C. Luthgrens and S. Moscovici. “Biased information search in group decision making.” Journal of Personality and Social Psychology, 78, 655-669. 2000

[TE11] UK Government technology strategy board (TSB) funded collaborative research project, see http://www.trust-economics.org/

[Tu84] B. W. Tuchman. “The march of folly: From Troy to Vietnam.” New York: Ballantine Books. 1984.

[Wa66] P. C. Wason. “Reasoning.” In B. Foss (Ed.) New horizons in psychology (pp. 135-151). Harmonsworth, Middlesex, England: Penguin. 1966

## Appendix 1: System architecture of the preference elicitation tool

The figure below provides a high-level view of the system architecture behind this tool. It is based on an engine that executes preference elicication workflows. Each step in the workflow can be configured in terms of the information that will be requested to the user and its graphical representation. The tool stores the gathered information into a centralised database, allowing for further post-processing and data mining.

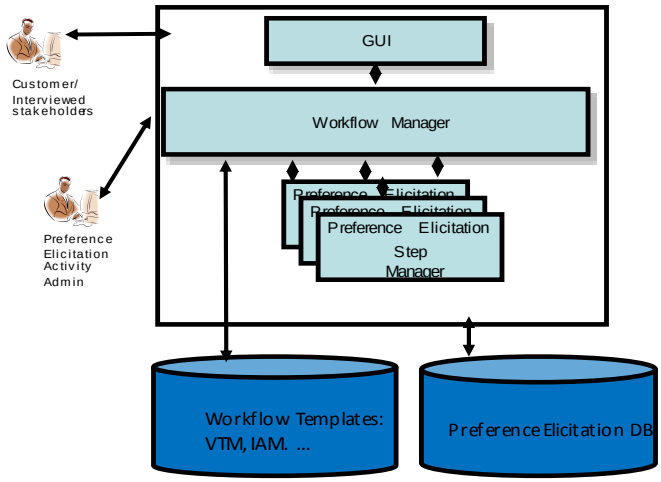


## Appendix 2: Summary of data analysis

| Phase | Result/Data Collected | Analysis |
|---|---|---|
| 3. Questions and Answers | 173 questions | Various, but main result was based on ratio of security related questions between the control and intervened groups |
| 6. Choice & Justification | 152 justifications | Various, but main results were based on ratio of security related justifications, and complexity of justifications between the control and intervened groups |
| 5a. Preference/Economic Framing | 6 ordered preferences over 4 components (table nn) | Participants preferences were compared with justifications |
| 6. Choice & Justification | 12 choices and 48 Likert scores (table nn) | Comparison between control and intervened groups |
| 7. Introspection | Judgments on the interventions | See discussion. |

## Appendix 3: Example/Illustration question and justifications

Example questions included:
Q1. What processes [do we have] to keep anti-malware up to date on clients
Q2. Do we have anything on the network that looks for unusual traffic, maybe an IPS?
Q3. Is training and [following of] procedures measured in any way?
Q4. To what regulations these customers [of our company] have to comply with, e.g. data protection legislation, etc.?
Q5. Precisely which countries do we operate in

Example justifications included:
J1. reduces ability for malware infections to gain admin rights on client systems
J2. reduces threat vector for zero day / early exploited vulnerabilities
J3. zero cost [and] no acceptance issues
J4. can impact on productivity
J5. high impact on user satisfaction and productivity [and] more expensive than HIPS option

## Appendix 4: Preferences expressed by the intervention group in phase 5a.

| Component | Intervened Subject 1 | Intervened Subject 2 | Intervened Subject 3 | Intervened Subject 4 | Intervened Subject 5 | Intervened Subject 6 |
|---|---|---|---|---|---|---|
| (Security) Risk | 1 | 1= | 1= | 1 | 1 | 3 |
| Compliance | 4 | 4 | 1= | 4 | 4 | 4 |
| Cost | 2 | 3 | 4 | 3 | 2= | 2 |
| Productivity | 3 | 1= | 3 | 2 | 2= | 1 |