Decision Support for Systems Security Investment

Yolanta Beresnevichiene, David Pym, Simon Shiu
Systems Security Lab
Hewlett-Packard Laboratories
Long Down Avenue, Stoke Gifford, Bristol BS34 8QZ, England, UK
yolanta.beres@hp.com, david.pym@hp.com, simon.shiu@hp.com

Abstract—**Information security managers with fixed budgets must invest in security measures to mitigate increasingly severe threats whilst maintaining the alignment of their systems with their organization's business objectives. The state of the art lacks a systematic methodology to support security investment decision-making. We describe a methodology that integrates methods from multi-attribute utility evaluation and mathematical systems modelling. We illustrate our approach using a collaborative case study with the security managers of a large organization divesting itself of its IT support services. The case study was validated against the experience and observations of the security managers and delivered, according to their judgement, useful results. Specifically, by integrating a mathematical model of system behaviour with an account of the utility of available security investment strategies, the case study has enabled them to understand better the trade-offs between the security performance and the operational consequences of their choices.**

Keywords: Information security; economics; systems modelling; decision support; risk management

I. INTRODUCTION

Individuals and organizations of all sizes and at all levels of criticality face increasingly severe and sustained threats to the security of their confidential information. An organization's CISO must determine an appropriate policy, process, and technological response to the threat faced by the organization 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.

The experience of security managers and researchers (e.g., [1, 3, 5, 11, 27]) suggests that accounting-based approaches to addressing this problem, employing return-on-investment-type calculations, cannot adequately address the operational and dynamic aspects of the analysis that is required. Indeed, very few business or IT stakeholders have a good understanding of how security choices affect business outcomes. Consequently, the role of CISO, for example, has grown from being a technical job with responsibility for IT security operations to being a senior business role that bridges the gap between business and information security strategy, see [2, 24]. The role is challenging for many reasons: the many stakeholders typically have different (often misaligned) incentives and preferences; there are complex inter-dependencies between investment choices; and there is a little useful information about future threats.

Recent work by the present authors and others (see, for example, [3, 5, 6, 13, 27]) has begun to develop a new methodology for addressing this problem that integrates two main approaches. On the one hand, we employ executable mathematical models of the underlying system captured within its dynamic threat and economic environments. On the other hand, we employ methods from economics — specifically, utility functions and their associated dynamic analysis — together with empirical data-collection techniques. Taken together, these technologies provide, we suggest, a valuable technique to support decision makers in addressing the information security investment question outlined above.

The methodology we employ is illustrated pictorially in Figure 1. We begin with a characterization of the problem, as presented by the decision-maker (e.g., the client organization's CISO). For example, the organization may be divesting itself of one of its constituent businesses and may wish to manage the change of status and access privileges of the affected staff. Associated with this divestiture, the CISO has a range of choices for the nature of the resulting system configuration (including security controls) and a range of preferences among the security outcomes. These preferences give rise to a formal expression of utility. The dynamics of this utility can then be explored by constructing an executable mathematical model of the system, in the context of its dynamic threat and economic environments. The construction of such a model must capture not only the preferences of the decision-maker, such as the CISO, in respect of the desired outcomes but also the architectural (and policy, and business process) constraints inherent in the problem.

Having constructed the model, its behaviour is simulated in the presence of a (stochastic) representation of the dynamic threat and economic environments — including, in particular, security investments — and its predictions are

validated against the preferences (expressed as a utility function) of the decision-maker (e.g., the CISO). The model may then be refined appropriately, as may the decision-maker's understanding of his preferences in response to the initial problem, which may itself be subject to reassessment and refinement.

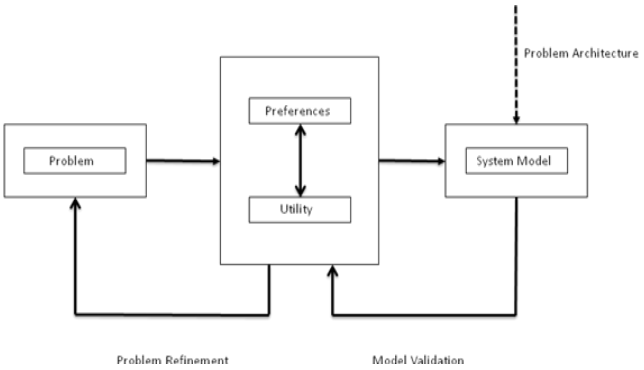

Figure 1: The methodology

The divestiture problem suggested above provides the basis of a case study, based on a collaborative project between the present authors and the security managers of a large financial services organization, which illustrates the deployment of our methodology in the context of a real-world problem. The primary research question is whether the methodology, possibly subject to refinement, can be applied usefully. If it can, then there is potential to completely change the way security decision-making is done. The secondary research question is to suggest solutions in the case study itself.

The structure of the remainder of this paper is guided by the methodology itself. Section II presents a more detailed account of the case study. Section III explains our use of utility functions as a basis for specifying the decision-maker's requirements. We also discuss the empirical methods employed in eliciting the decision-maker's preferences, and the relationship between the utility function and the dynamics of the system model. Section IV introduces in more detail both the mathematical modelling approach — beginning with its conceptual basis, with a very brief discussion of the implemented tool — and the specific model employed in the case study, illustrated pictorially. Section V considers the validation step, mapping the results of the simulations back to the decision-maker's preferences. Finally, in Section VI, we discuss what conclusions can be drawn from this early-stage exploration of integrating economic and mathematical systems modelling in an empirically based setting.

II. The Case Study

To explore the methodology we have outlined and, in particular, to address a real challenge facing the security managers in the financial services organization, we performed a case study of the ongoing de-perimeterization of their organization; see, for example, [25]. Within the de-perimeterization project, a typical example involves the divestiture of a business function or service, so that a business or service that initially existed entirely within the enterprise firewall, and which involved contracted employees, switches to being operated by third party employees accessing applications from outside the firewall. In such situations, the information security concern relates to the increased risk of breaches that may be introduced by more relaxed network access arrangements, changes in personnel culture, and changing contractual agreements.

Various security mechanisms can be considered to control communications between users (and their associated endpoints) and servers/applications. Such mechanisms cost money and can adversely affect the user or business process. In general, the problem is to determine which security portfolio will, at appropriate cost, provide the best trade-off between reducing risks and maintaining business priorities. The security controls that are often considered during the de-perimeterization of part or the whole of organization's network include the following: some type of virtual desktop environment with different restrictions and monitoring; controls enforcing stronger authentication for direct access (especially for servers and applications that cannot be moved to be accessible via the virtual desktop environment); intrusion detection systems (IDSs) to monitor and alert based on inappropriate network activity; regular access and privilege review to ensure there is no creep up of the number of users with multiple privileges. Different combinations of these controls will have different effects on different aspects of the system and its managers' confidence in its status, such as the likelihood or impact of certain types of breaches, the level of assurance/knowledge that breaches are detected, the performance of the business process, the costs of running the IT systems, or the security investment costs. For example, certain restrictions on the virtual desktop might be better at preventing malicious confidentiality breaches, whereas others will apply to inadvertent availability problems. Some mechanisms will make it more difficult for staff to inadvertently or maliciously cause breaches, but might also slow down their productivity. Alternatively, IDSs and other forms of monitoring might affect system latency, but improve awareness of the threat situation and so improve assurance.

It is important to emphasize that the study we describe is applied work, intended to deliver practical advice to managers (our collaborators) facing real operational problems. Given this context, our solution is necessarily approximate and incomplete, relying in some aspects on the experience and judgment of the security managers, elicited by iterations of the modelling methodology, and the modellers.

III. Trade-offs and Utility Functions

Once the decision-maker has adequately characterized the problem, with a range of (competing) attributes and objectives identified, it is necessary to determine to what extent the

objectives must be achieved for a solution to the problem to be acceptable; that is, we must determine the decision-maker's preferences for acceptable trade-offs between the various attributes and express them in a quantifiable form.

We adopt standard techniques from economics [17], as described in the systems modelling context in [3, 12], and employ utility functions of the form

$$U = \omega_1 f_1(C - \hat{C}) + \omega_2 f_2(A - \hat{A}) + \omega_3 f_3(I - \hat{I}),$$

where C, A, and I represent the outcomes — here, for example, confidentiality, availability, and investment — we care about, and $\hat{C}$, $\hat{A}$, and $\hat{I}$ represent the decision-maker's targets for these outcomes. The functions f_i $(1 \leq i \leq 3)$ represent the decision-maker's tolerance for variance from the targets. The weights ω_i $(1 \leq i \leq 3)$ represent the decision-maker's preferences between the component outcomes. Of course, the utility function may have many components.

In the simplest case, we set the f_is to be quadratic functions. This choice, which has a well-supported theoretical basis [17], captures diminishing marginal utility and implies, since quadratics are symmetric about their maxima, that the decision-maker is equally tolerant for going over or under target. For example, if the outcome component is cost, overspending by £100 is just as bad as under spending by £100. In most practical situations, however, the decision-maker's preference will be asymmetric and it is necessary to use functional forms such as Linex functions [26, 28, 13], of the form $f(x) = (e^{\alpha x} - \alpha x - 1) / \alpha^2$, which capture this asymmetry appropriately (α is a parameter).

Having established the form of the utility function, we consider its expected value as the components vary over time. The dynamic models employed in economics (for a security example, see [13]) employ a set of system equations that describe the dynamics of the components in the presence of stochastic shocks. Instead of a set of equations, we employ a mathematical system model which captures the structure of the system in terms of its key components (see Section 4) and which can be executed in order to simulate the behaviour of the system in the presence of stochastic shocks. The structure of such a model allows the (expected) values of the components of the utility function to be calculated.

In the case of any particular model, such as that developed here, the components of a utility of interest must be identified. In our case study, these were identified via a process of multiple iterations with the decision-makers. The process sought to determine their primary concerns, the changes they expected over the future years of interest, their investment options, and the expected consequences of these investments, including the preferences between outcomes associated with each investment option. The process was implemented using structured discussion within which the consequences of focusing on certain components were considered.

Initially, the components considered included cost, confidentiality, and availability. These attributes clearly trade off against one another, as each (confidentiality) mechanism that restricts or reduces access naturally makes the system less available, and vice versa. As different types of availability outcomes were discussed, it became apparent that the real issue was the effect on the business function (as opposed to system or network uptime, bandwidth, or latency). Similarly, confidentiality shifted to cover many forms of breaches, including the integrity of transactions, data leakage, and unauthorized or even unaccountable system activity. It was clear that for each of these there was a desire to reduce the number of breaches, but also to know (and communicate) the effectiveness of the methods of reducing breaches.

As a result of this empirical work, the utility components for the case study became *breach prevention*, *assurance*, and *business performance* — corresponding conceptually to the security components (such as C, A, etc.) in the utility expression above — and cost — corresponding to I in the utility expression above. In the case study, the business performance component of utility represents the performance of IT support staff in response to support–job requests. The next step was to elicit the tolerance for how much should be achieved in each of these components. For example, in order to elicit the decision-maker's preferences for the form (e.g., asymmetry, gradient) of marginal utility either side of target, the use of both quadratic and Linex formulations of the dependencies of the utility function on components were explored using a structured questionnaire.

As experience would suggest, the decision-maker's utility function in this case study is, to varying degrees, asymmetric in all of its components. For example, the marginal utility of breach prevention has steeper gradient below target than above. The assignment of form employed in this study is imprecise. Nevertheless, we were able to use this information to inform the design of the next questionnaire, used to elicit the decision-maker's preferences between outcomes.

A. Capturing the Decision-maker's Preferences

We have explained that, in the case study, the decision-maker's desired collection of attributes — essentially, this is the problem characterization — were elicited via structured discussion. It was also necessary to elicit the decision-maker's preferences between outcomes. To this end, we focussed on single measures that would represent each utility component, and presented each outcome as a 4-tuple, consisting of proportion of breaches against overall access activity (i.e., breach prevention), proportion of detected breaches against overall breaches, (i.e., assurance), proportion of SLA violations against overall job-requests (i.e., business performance), and cost. From these, we created simple preference questionnaires each consisting of around 100 value pairs related to the components in the 4-tuple. For example, value pairs were created for breaches and SLA violations, where values for the proportion of breaches against accesses ranged from 0.15 to 0.01 with different values for proportion

of SLA violations. The decision-maker had to evaluate and rate each pair within the scale of 1–6, where 3 would represent an acceptable outcome, 6 would be strongly unacceptable, and 1 would be a highly desirable outcome.

Figure 2 plots the decision-maker's preference values against breaches and business performance and Figure 3 plots preference values against the proportion of accesses that are breaches and the proportion of breaches that are detected. The prevalence of diamonds and squares towards the bottom left indicates the preference for as few breaches as possible, a high detection rate, and some tolerance for SLA violations.

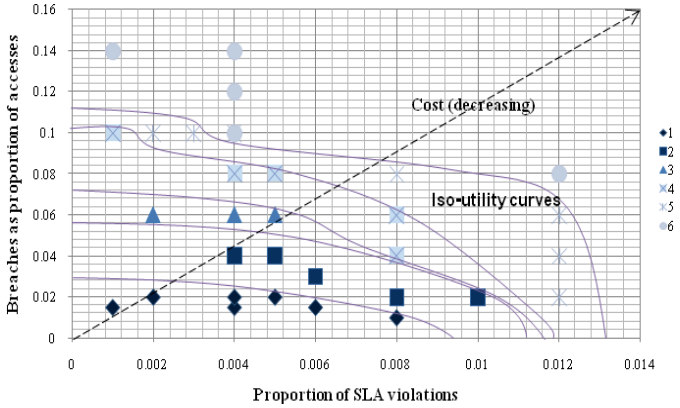

Figure 2: Decision-maker's preferences against breaches and SLA violations

The questionnaire used allowed us to establish where there was high intolerance for certain outcomes and the broad preference relationships between breaches and performance, and between breaches and assurance.

As more empirical data was obtained, iterated uses of the questionnaire proved more effective, and the iso-utility curves (the asymmetry in the utility is discussed briefly in Section 5.1) shown in Figures 2 and 3 provided a better connection with the security manager's experience. Cost can be overlaid on these preference values and utility curves, as indicated by the arrows of decreasing cost. This data was used to relate predicted outcomes (from the executable model) to the managers' initial preferences, allowing refinement. Given more time with the security managers, the next step would have been to run a further questionnaire to identify more accurately the positions of the iso-utility contours.

We emphasize that the empirical studies (with all their well-documented attendant difficulties [12, 14, 15, 18]) are beyond the scope of the present paper. For now, we require a plausible way to proceed to test the feasibility and value of the overall framing and modelling methods in the decision process, deferring consideration of more rigorous preference-elicitation methodologies to another occasion.

As we have explained in the introduction, the elicited preferences are used to condition an executable, mathematical system model, described in some detail in Section IV.

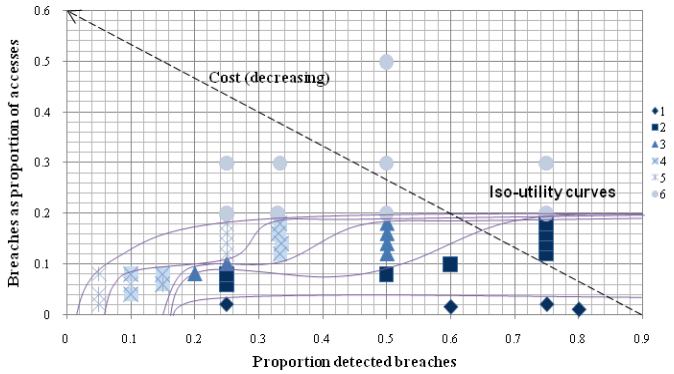

Figure 3: Decision-maker's preferences against overall breaches and their detection rate

IV. SYSTEMS MODELS AND SIMULATIONS

We have described our whole-systems view of the security decision-making problem. A key component of that view is our mathematical modelling of the underlying architecture and processes. Our approach, the mathematical basis of which is presented in [5, 8, 9, 29, 30], is grounded firmly in mathematical logic, computation theory, and probability theory, but employs well-developed, implemented tools.

Our approach views a system as having the following key conceptual components [4, 8, 9, 29, 30]: **Environment**: All systems exist within an external environment. We may seek to model the structure of the environment, in which case we treat the environment as a system of interest in itself; typically, however, we treat the environment as a source of events that are incident upon the system of interest according to given probability distributions; **Location**: The components (i.e., resources; see below) of a system of interest are typically, distributed around a collection of places. Different places are connected by oriented links; **Resource**: The notion of resource captures the components of the system that are manipulated by its processes (see below). Resources include things like the components used by a production line, the system operating staff, and money; **Process**: The notion of process captures the (operational) dynamics of the system. Processes manipulate resources in order to deliver the system's intended services.

This framework, and the Gnosis modeling tool, are described in detail in [6, 8, 9, 29, 30], with related work in, for example, [21, 22, 23].

A. The Model Employed in the Case Study

The system model created for the case study represents the access activity of IT support staff in the de-perimeterized network environment and explores the outcomes for 4-tuple measurements as described in Section 3.1. These measurements are gathered through several simulations of the model, each under different combinations of the security controls as listed in Section II.

Specifically, the model captures the process of IT staff responding to support-job requests and accessing numerous

internal systems, via various access protocols. We assumed, guided by the experience of the managers, that untrustworthy staff would take opportunities to engage in unauthorized activities (including harmless, justifiable accesses, over-zealous trawling, or significant breaches) in addition to the legitimate job-related activities. Depending on the access protocol used, we also assumed, based on the experience and observations of the managers, certain success rate for a breach and its detection by monitoring controls. Any additional security controls introduced would either reduce the likelihood of a breach occurring (this would be with restrictions on virtual desktop access, and direct access controls), or improve would detection rate (any monitoring controls). At the same time, the additional controls put some burden on support staff, thus decreasing their job turnover rate. Additional access controls, for example, often require extra authentication. Also, these controls are usually centralized, thus requiring a central server to be always online. If it fails, the system becomes inaccessible for a number of hours. Figure 4 shows the general structure of the model. It consists of two main parts. One part models the activities of the IT support staff, mainly the job-request processing. This task requires multiple accesses to the systems either though a virtual desktop or direct access protocols. The other part evaluates each access and determines the likelihood of its resulting in unauthorized activity and breaches.

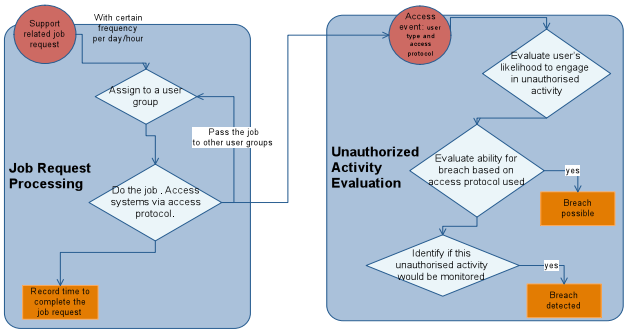

Figure 4: General structure of the system model

The diagrams should be interpreted as follows: the circular components represent events incident upon the system from the environment; the rectangular endpoints correspond to the resource components of the model, and provide the measurable quantities for utility calculations (we make no use of location in this model); finally, the process dynamics of the model is captured by the arrows connecting events to resources via key computation steps, denoted by diamonds.

The model requires some initial assumptions to be made about the initial state of access requirements, IT staff trustworthiness, and the general job-request frequency and turnover rate. Table 1 below summarizes these assumptions, which were based on the experience and observations of the IT operations and security managers in the organization and which elicited via multiple iterations of model-execution and model-refinement.

The job-request processing part of the model schedules new jobs every hour and assigns them to IT staff.

Table 1: Initial state assumptions for the model

For job-request processing
Support job frequency: 1 every hour
Time1 taken to do the job (1st user/pass): between 2h and 5h
Time2 taken to do the job (2nd and further users/passes): between 0.5 and 1 hours
Multi-group ratio (users in more than one group): 0.5 (i.e., 50% of users are in multiple groups)
Job redo ratio (more than one user works on it): 0.01/multi_group_ratio
Standard SLA required job processing time: 6 hours
For unauthorized activity evaluation
User trustworthiness: 90% trusted
Non-job-related access ratio: 0.005
Access protocol ratio: 55% of accesses go through virtual desktop, 45% are direct accesses
Ability to engage in unauthorized activity is at 0.4 via virtual desktop access with 0.7 detection rate (when monitoring is in place), and at 0.75 via direct access with 0.4 detection rate

Each IT staff (corresponding to a system user) accesses the systems to work on the job using an access protocol selected based on the access protocol ratio in Table 1. Also, occasionally, non-job-related access is triggered, corresponding to 0.5% of overall accesses (ratio 0.005 in Table 1).

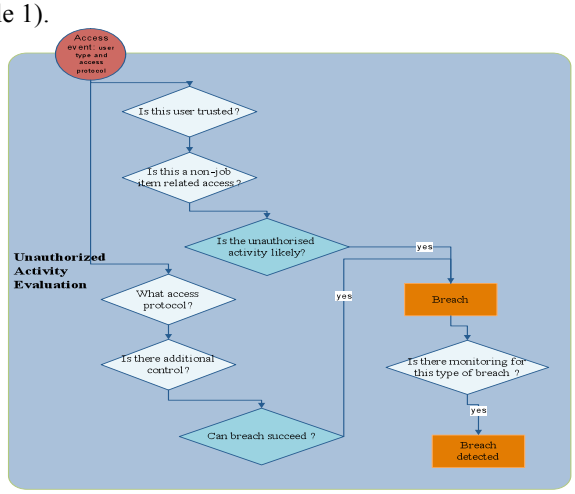

Figure 5: Detailed model for evaluation of unauthorized activity

Depending on to how many groups the user belongs, the job could be passed to another user after certain time (Time1 in Table 1). The second and any subsequent users take additional time (Time2) to finish the job. In the end we arrive at the measure of the overall time taken to complete a support job request. If it exceeds the SLA-dictated time, the task is registered as an SLA violation. The unauthorized activity evaluation part of the model is used to determine the likelihood of a user engaging in unauthorized activity and the ability of this user to successfully execute a breach. A detailed breakdown of this part of the model is shown in Figure 5

Based on the advice of the managers, we assume, initially, that the overall likelihood for a user to engage in unauthorized activity is at 0.001. This increases 10-fold if the user accessing the systems is not trustworthy (based on the ratio in Table 1). It would double if the access were determined not to be job-relevant. Depending on the access protocol employed, the

ability for the user to successfully execute a breach differs. We assume that breaches are more likely to succeed via direct access to systems (with probability of 0.75) than via virtual desktop access (with probability 0.4).

This is, of course, a very simplified view on how breaches might arise. A more rigorous analysis, based on the attacks trees [10, 19] and hacker behaviour [16] could be used to arrive at more grounded and realistic probabilities regarding the breach success rate. In this research, however, we have focussed not on analyzing the internal/external attacker behaviour, but rather on making reasonable assumptions. Assumptions were also made about the breach-detection rate related to each protocol (as in Table 1). These rates are used to determine the proportion of detected breaches.

A. Impact of control investment choices

The investment choices made on additional security controls will have impact both on job request processing and on unauthorized activity parts of the model. As can be recalled from Section II, four additional security controls were considered in the case study: centralized access controls (authentication plus enforcement) and monitoring for virtual desktop environment (VD) controls, additional authentication when using direct access (DA) controls, network-level intrusion detection monitoring (IDS), and regular privilege (Priv) review. Based on iterations with IT operations and security teams, we decided on a changed set of assumptions about the effect each of these controls would have on the dynamics in the model, summarized in Table 2.

Table 2: Effect of additional controls on the assumptions

For job-request processing
VD controls would be more prone to failure: failure likelihood at 0.05 every 2 weeks and the associated down time at 5 hours on average
DA controls would lengthen the job processing time by 5-15 minutes
IDS have no impact on job processing time
Priv. review would decrease the group ratio from 50% users in multiple groups to 10% in multiple groups. This would result in large job redo ratio.
For unauthorized activity evaluation
VD controls would reduce the ability for breach from 0.4 to 0.15 and increase detection rate to 0.85
DA controls would reduce the ability for breach from 0.75 to 0.4.
IDS would increase the detection rate of breaches via direct access to 0.6
Priv. review would decrease the likelihood of non-job-related access activity from 0.005 to 0.001, thus reducing the likelihood of user engaging in an unauthorized activity.

B. Results from Simulations

The simulations of the model were run over a one-year period, 100 times, and were performed for every combination of the control investment portfolio, starting from one of the four controls being turned on, then two, and so on until all four controls are turned on. During the simulations we gathered the measures corresponding to the 4-tuple components as described in Section II.

The results in Figure 6 show the proportion of breaches and SLA violations measured across 8 control combinations. The first control combination in the chart, when VD, DA controls and privilege review are off, could be the starting position in the organization that has just basic controls. The results indicate that the highest reduction in breaches is achieved when virtual desktop controls and direct access controls are implemented and when privilege review decreases the multi group ratio to 0.1. This investment option, however, has one of largest impact on staff productivity by increasing the proportion of SLA violations by 0.0108.

Next, the IDS monitoring control was considered, in addition to the previous controls. Figure 7 shows results comparing effect on breaches and their detection rate under similar control investment options as before, but with IDS turned on for the last four cases (the privilege review was not considered in this case as it has no effect on detection rate). As can be seen from this chart, the detection rate increases the most when all three controls are turned on.

These results highlight that all four of the security controls under consideration must be implemented to achieve the biggest security benefits in terms of reduction of breaches and detection rate. This outcome comes is unsurprising, but the results highlight other security options that achieve results close to this best option. For example, with both VD and DA controls, but with no privilege review, a very similar reduction in breaches is achieved, differing from the best option only by 0.0008. This option increases the SLA violation rate only very slightly compared to the best option. Overall, privilege review appears to be most likely to increase the violation of SLAs. It should be noted that these results are based on the assumption — informed by the experience of the security managers — of a fairly trustworthy IT staff population: 90% are trusted. If the trusted population is reduced to 60%, the same simulations indicate that the level of unauthorized activity more than doubles. In such cases, other controls might become necessary to keep unauthorized activity and breaches at acceptable levels.

V. VALIDATION: MAPPING THE SIMULATION RESULTS TO THE DECISION-MAKERS PREFERENCES AND UTILITY FUNCTION

We now examine how the results based of the simulations compare with the decision-maker's preferences for the same measurement pairs of breaches, SLA violations, and breach detection rate. This helps us determine which of the (eight) investment options result in outcomes that match closest to the highest rankings given by the decision-maker.

We position the results from Figures 6 and 7 alongside the previously extracted decision-maker's ratings. Figure 8 shows results from Figure 6 alongside the preference values from Figure 2 for breaches and SLA violations. The chart indicates that at least 4 control choices demonstrate outcomes that fall within the area of the highest rated preferences.

These include investment options where privilege review is turned off, with VD and DA controls either off or on. There is a starting position in which all controls are off.

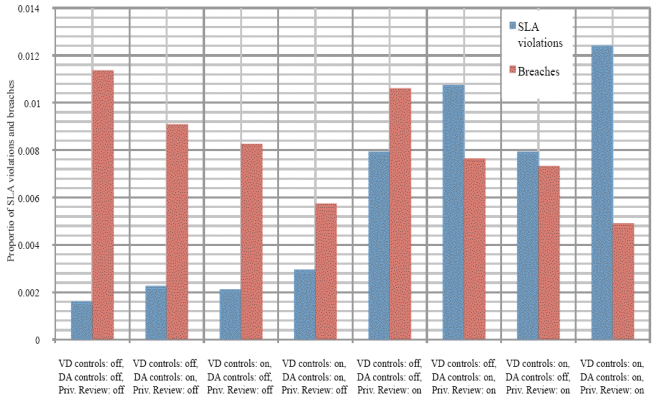

Figure 6: Comparison of breaches and SLA violations

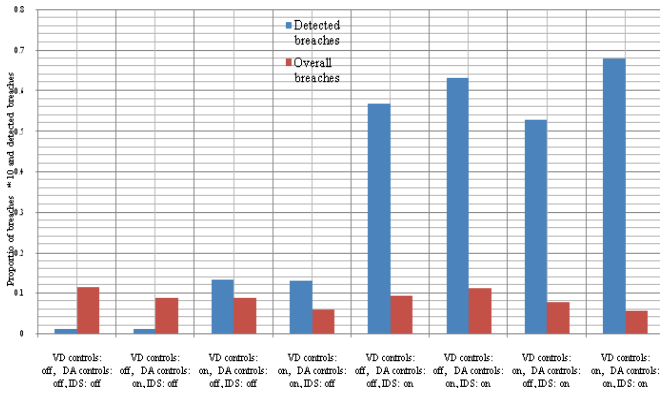

Figure 7: Comparison of breaches and detection rate (privilege review is off)

This means that the interviewed decision-makers perceive that the organization is currently at an acceptable level in terms of breaches and SLA violations, but we need to recall that the results here are based on an assumption of fairly trustworthy IT staff population (90% are trusted). Since the decision-maker believes that trustworthiness of IT staff, and the associated threat-environment, will deteriorate in the wake of the divestiture, the organization must reduce the growing breach level. The chart points suggests the investment option where both VD and DA controls are implemented as achieving the lowest level of breaches with an acceptable level for SLA violations.

Figure 9 positions results from Figure 7 alongside ratings from Figure 3 in order to compare preferences against the measurements of overall breaches and their detection rate. Once again, four options fall within the area of highest ranked preferences, all being where IDS is on. The best within those would probably be the one with the highest detection rate.,for which both VD and DA controls are on together with IDS .

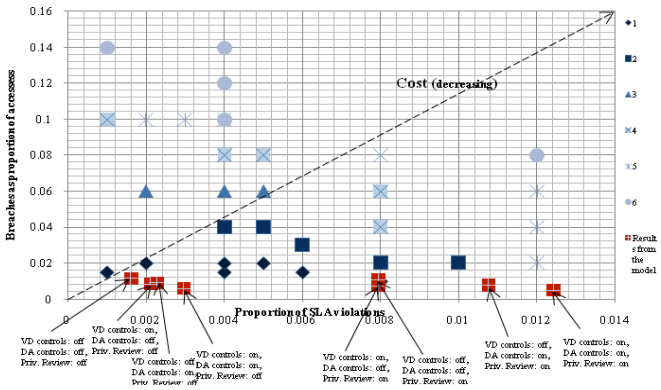

Figure 8: Simulation results mapped to decision-maker's preferences (Figure 2)

Overall, it seems that of the eight security investment options that were experimented with during simulations, the decision-maker would most prefer the option with VD and DA controls on, IDS on, but privilege review off. Other options, with either VD or DA being on or off, are also candidates.

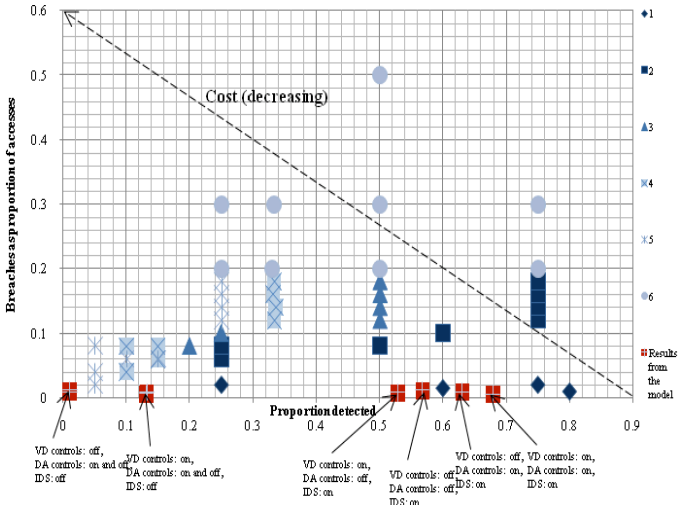

Figure 9: Simulation results mapped to the decision-maker's preferences (Figure 3)

A. Cost and Overall Utility

The ratings and graphs do not reflect preferences on cost. Figures 2 and 8 show, unsurprisingly, a preference for low levels of SLA violations and breaches. Achieving these outcomes will be expensive (cost is shown to be decreasing from the bottom left). Similarly, a small number of breaches that are all observed is desirable and expensive (cost is highest in the bottom right of Figures 3 and 9). So the cost information in the graphs for each portfolio is a necessary component if conclusions are to be drawn about which investment options provide the most appropriate trade-off between unauthorized activity, SLA violations, observed breaches, and cost.

Using a utility function of a particular form constitutes a framing of the problem. For this case study, the ratings, and the iso-utility curves, implicitly carry information about the form of the utility function, though translating this into formal

parameters is beyond our present scope. The point is that the form of utility function and the initial intuitions about each of the components together provide a traceable and coherent explanation of the data. For this reason, the acknowledged framing (i.e., chosen form of utility function) adds value.

VI. CONCLUSIONS

The case study demonstrates that it is feasible for a security team, unfamiliar with economics or systems modelling, to work through this new methodology. Broadly, the process was perceived by the managers to be objective and rigorous. More specifically, although the team recognized the crudeness of some of the assumptions, and of the elicitation of preferences, they were comfortable to proceed, see the implications and re-discuss the initial predicates. The specific results aligned with intuitions, and confirmed the broad strategy being followed. The process was seen as useful both for the team's confidence in its choices and for providing better grounded and more transparent due diligence to other (e.g., non-security) stakeholders. We interpret this feedback on the case study as indicating that there is pragmatic value in integrating models of utility from economics with the executable mathematical modelling approach in the development of tools to support investment decision-making in information systems security.

Following further theoretical work on the methodology, including examining appropriate preference-elicitation mechanisms and handling imprecision [15], an interesting next case study involve stakeholders not part of the security team. This would allow consideration of the effect of the decision-maker's choices on a utility for the broader organization.

ACKNOWLEDGEMENT

We are grateful to Max Heitman and Peter Ventura at Citigroup for their detailed and considered engagement with our case study.

REFERENCES

[1] Anderson, R.: Why information security is hard — an economic perspective. *Proc. 17th Annual Computer Security Applications Conference*, 358–365, New Orleans: IEEE, 2001.

[2] Ashenden, D. Information Security Management: "A Human Challenge?" Information Security Technical Report (2008), doi: 10.1016/j.istr.2008.10.006.

[3] Beautement, A., Coles, R., Griffin, J., Ioannidis, C., Monahan, B., Pym, D., Sasse, A., Wonham, M.: Modelling the Human and Technological Costs and Benefits of USB Memory Stick Security. In: Managing Information Risk and the Economics of Security, M. Eric Johnson (editor), 141–163. Springer, 2009.

[4] Birtwistle, G.: *Discrete event modelling on Simula.* Springer, 1987.

[5] Beres, Y., Griffin, J., Shiu, S., Heitman, M., Markle, D., Ventura, P.: Analysing the Performance of Security Solutions to Reduce Vulnerability Exposure Windows, *Proc. Annual Computer Security Applications Conference*, 33–42, Anaheim, California: IEEE, 2008.

[6] Collinson, M., Monahan, B., Pym, D.: *A Discipline of Mathematical Systems Modelling*. College Publications, London, 2010 (forthcoming).

[7] Collinson, M., Griffin, J., Monahan, B.,: Towards a Toolset Framework for Gnosis. Forthcoming HP Labs Technical Report, 2010.

[8] Collinson, M., Monahan, B., Pym, D.: A Logical and Computational Theory of Located Resource. *Journal of Logic and Computation*, 19(6): 1207–1244, 2009. doi:10.1093/logcom/exp021.

[9] Collinson, M., Pym, D.: Algebra and logic for resource-based systems modelling. *Math. Struct. Comp. Sci*. 19:959-1027, 2009.

[10] Dawkins, J., Hale, J.: A systematic approach to multi-stage network attack analysis. *Proceedings of Second IEEE Int. Information Assurance Workshop*, 48–7, April 2004.

[11] Gordon, L., Loeb, M.: *Managing Cybersecurity Resources*. McGraw Hill, 2006.

[12] Hersey, J.C., Kunreuther, H.C., Shoemaker, P.J.: Sources of bias in assessment procedures for utility functions. *Management Science* 28, 936–953, 1982.

[13] Ioannidis, C., Pym, D., Williams, J.: Investments and trade-offs in the economics of information security. *Proc. Financial Cryptography and Data Security*, Dingledine, R. and Golle, P. (editors), LNCS 5628: 148—166, Springer, 2009.

[14] Jaffrey, J.Y.: Some experimental findings on decision-making under risk and their implications. *European Journal of Operational Research* 38, 301–306, 1989.

[15] Jimenéz, A., Ríos-Insua, S. Mateos, A.: A decision support system for multi-attribute utility evaluation based on imprecise assignments. *Dec. Supp. Syst*. 36, 65–79, 2003.

[16] Jonsson, E., Olovsson, A.: Quantitative Model of the Security Intrusion Process Based on Attacker Behaviour. *IEEE Transactions on Software Engineering* 23(4), 235—245, 1997.

[17] Keeney, R.L, Raiffa, H.: Decisions with Multiple Objectives: Preferences and Value Trade-offs. Wiley, New York, 1976.

[18] McCord, M., de Neufville, R.: Lottery equivalents: reduction of the certainty effect problem in utility assessment. *Management Science* 32, 56–61, 1986.

[19] Moore, A., Ellison, R., Linger, R.: Attack Modeling for Information Security and Survivability. Carnegie Mellon Software Engineering Institute. Technical Report CMU/SEI-2001-TN-001, 2001.

[20] Nicol, D., Sanders, W., Rivedi, K.: Model-based evaluation: from dependability to security. IEEE Transactions on Dependable and Secure Computing 1(1), 48–65, 2004.

[21] Pidd, M.: Tools for Thinking: Modelling in Management Science, Wiley, 2003.

[22] Pidd, M.: Complementarity in systems modelling. In: Systems Modelling: Theory and Practice (M. Pidd, editor), Wiley, 2004.

[23] Pidd, M.: Computer Simulation in Management Science. Fifth Edition, Wiley, 2004.

[24] Rasmussen M, Stamp P.: Where Security Reports Reflects Expanded Role And Responsibilities. Forrester 2005.

[25] Stamp, P.: Jericho Forum Looks To Bring Network Walls Tumbling Down. Forrester, 8 July, 2005.

[26] Varian, H.: A bayesian approach to real estate management. In S.E. Feinberg and A. Zellner, editors, *Studies in Bayesian Economics in Honour of L.J. Savage*, 195–208, North Holland, 1974.

[27] Yearworth, M., Monahan, B., Pym, D.: Predictive Modelling for Security Operations Economics. Presented at the Workshop on the Economics of Securing the Information Infrastructure (WESII), 23—24 October, 2006, Washington DC. Available as HP Labs Technical Report HPL-2006-125.

[28] Zellner, A.: Bayesian prediction and estimation using asymmetric loss functions. *J. American Statistical Association* 81:446–451, 1986.

[29] Collinson, M., Monahan, B., and Pym, D.: Semantics for Structured Systems Modelling and Simulation. In *Proc. Simutools 2010*. ICST: ACM Digital Library and EU Digital Library, 2010. ISBN: 78-963-9799-87-5.

[30] Gnosis. http://www.hpl.hp.com/research/systems_security/gnosis.html.