Investments and Trade-offs in the Economics of Information Security

Christos Ioannidis[1], David Pym[2⋆], and Julian Williams[3]

[1] School of Management, University of Bath
Bath BA2 7AY
England, U.K.
c.ioannidis@bath.ac.uk
[2] Hewlett-Packard Laboratories
Bristol BS34 8QZ
England, U.K.
david.pym@hp.com
[3] University of Aberdeen Business School
Aberdeen AB24 3QY
Scotland, U.K.
julian.williams@abdn.ac.uk

Abstract. We develop and simulate a dynamic model of investment in information security. The model is based on the recognition that both IT managers and users appreciate the trade-off between two of the fundamental characteristics of information security, namely confidentiality and availability. The model's parameters can be clustered in a manner that allows us to categorize and compare the responses to shocks of various types of organizations. We derive the system's stability conditions and find that they admit a wide choice of parameters. We examine the system's responses to the same shock in confidentiality under different parameter constellations that correspond to various types of organizations. Our analysis illustrates that the response to investments in information security will be uniform in neither size nor time evolution.

1 Introduction

Information security and network integrity are issues of the utmost importance to both users and managers. The cost of security breaches and fraud is considerable and Anderson et al (2007) [1] provide a comprehensive review of the issues both technical and legal and offer a set of very useful recommendations. Such issues constitute growing concerns for policy makers, in addition to the legitimate concerns of the specialist technological community of experts. As the importance of networks increases for all individuals who act as both providers and consumers of information, the integrity of such systems is crucial to their welfare. In the presence of threats to the system, agents must decide the amount of resources required to maintain the system at acceptable operational states.

⋆ Also: University of Bath, Bath BA2 7AY, England, U.K.; d.j.pym@bath.ac.uk

Finding solutions to this resource allocation problem is therefore an important part of the work of IT managers. As with all such decisions, expenditure in protecting a system has an opportunity cost because resources can be deployed for other useful purposes, a situation that requires the manager to demonstrate the desirability of such expenditure given an objective that takes into account that such protection costs are fully justified in the light of a well-specified objective.

The calculation of the optimal investment in information security given the system's configuration is a subject that is relatively recent as the research literature has and focused almost exclusively on technological solutions without recourse to the associated financial costs and the behavioural changes required to implement such purely technological solutions. The economics of information security within the context of an optimizing framework has been addressed relatively recently by Gordon and Loeb (2002) [4], who provide an extensive list of references that address technological issues in information security and point out the distinct lack of rigorous economic analysis of the problem of resource allocation in information security. Gordon and Loeb adopt a static optimization model where IT managers calculate the optimal ratio of investment in information security to the value of the expected loss under different assumptions regarding the stochastic process that generates the security threats. Within the framework of the model, we conclude that a risk-neutral firm should spend on information security just below 37% of the value of the expected loss that will occur in the event of breach.

The model relies on rather restrictive assumptions and has prompted lively debate regarding the 'optimal' ratio of investment in information security. What is of interest is that the relationhip between investment in information security and vulnerability is not always a monotonic function. Hausken (2006) [6] by postulating an alternative functional form of vulnerability shows that the ratio cannot be supported. In similar vein, Willemson (2007) [9] introduces the notion of the existence of a level of expenditure of information security that removes all threats, as an additional parameter, thus completely securing the information. Under this specification the 'optimal' ratio can vary according to the value of this parameter. The author constructs examples where optimal investment ranges between 50% and 100% of the value of information that is protected.

All such models share a number of characteristics such as the knowledge of the 'monetary' value of information that is safeguarded and in addition the very metric of infromation security as such is not defined. It is simply stated in its 'negative' appearance as the value of the loss. Gordon and Loeb concede that the constituent components of the composite 'service' of information security may not be mutually consistent but given the requirements of their model and the assumption that all information can be valued by such decomposition is not necessary for the analysis undertaken.

In this paper, we develop a dynamic model that acknowledges the existence of trade-offs between the fundamental characteristics of information security, namely confidentiality, integrity, and availability (for simplicity here, we restrict to confidentiality and availability; cf. Beautement et al. (2008) [2]). Our inspi-

ration, and justification, for this analysis is derived from an empirical study Beautement et al. (2008) [2]. Specially, two of us (with others), studied and analyzed the costs and benefits of USB memory stick encryption in the context of the use of USB memory sticks by the staff of an investment bank.

The analysis of the paper Beautement et al. (2008) [2] can be summarized conveniently as follows:

- We observe that, for very well-motivated business reasons, the staff of an investment bank use USB memory sticks to store and transfer information at and between a variety of different locations with differing threat and security profiles;
- We observe, and collect supporting empirical data to the effect, that there are availability-driven incentives not to deploy technologies that promote confidentiality — essentially, it is highly inconvenient, and embarrassing, for the bankers to be unable to remember the necessary password in the presence of the client, and may lead to loss of business;
- We build executable mathematical models of the lifecycle of a USB stick which allow the exploration of the influences of various forms of investment — in training, IT support, and monitoring — on the use of encryption for USB memory sticks;
- We observe that the behaviour of these models does indeed support the existence of a trade-off between confidentiality and availability in this context.

Of course, technological solutions, such as biometric access control, may largely solve this particular problem, but we suggest that the methods and models that we are developing will be of use in a wide variety of situations.

Note that, for the purposes both of the study described above and of the model prsented in this paper, we are concerned with the following notions of confidentiality and availability:

- We consider the confidentiality of the system as represented by the extent to which the system is protected against unintended exposures of information. To this extent, to do not consider the confidentiality of information exposed by given breach; rather, we are concerned with the extent to which is protected against further breaches;
- For simplification, we neglect integrity in the model presented herein. In the context of the study of Beautement et al. (2008) [2], corruption of data as a consequence of the use of USB memory is a relatively minor issue, and the model we present herein should be considered to be potentially applicable only in situations in which such a simplification is justifiable. Clearly, other simplifcations are possible and may be supported by different circumstances and examples. We defer a more comprehensive discussion of the variety of models supported by the general framework introduced in § 2, within which integrity can be incorporated, to another occasion;
- Again, as a simplification, we adopt a simple proxy for availability: the degree of inter-connectedness of the system's components, which may be thought of as a measure of the size of the 'attack surface'.

Managers optimize well-defined objective functions in terms of such elements and recognize that investment is costly. The system's state equations determine the system's operational efficacy and the managers' optimal responses when under 'attack' are defined by altering the system's inter-connectness and by the acquiring new investment in information security. All the parameters have explicit behavioural and technical interpretations and allow for the classification of managers' behaviour and the system's architecture.

In § 2, we begin with a summary of the simplifying assumptions, motivated by the study presented in Beautement et al. (2008) [2], employed in this paper. We also provide, following Gianni and Woodford (2002) [3], a brief summary of the general linear stabilization problem and its solutions, and discuss briefly its use, by Nobay and Peel (2003) [8], in monetary policy with asymmetric preferences. In § 3, we describe our model in detail, providing the necessary system of differential and integral equations, together with their interpretation in terms of the concepts of information security. In § 4, we provide a range of examples of constellations of the model's parameters, corresponding to organizations with contrasting information security preferences and management policies, and provide graphs of simulations illustrating the impulse–response of these organizations to a single (exogenous) unit-shock to confidentiality. Finally, in § 5, we provide a range of observations, variations, and extensions about our modelling framework. We provide also two appendices: the first explains the discretization of our models used to generate our simulations; the second explains how our quadratic form of loss functions derives from basic concepts of utility theory.

2 CIA, Investments, and Trade-offs

Organizations deploy systems technologies in order to achieve their business objectives. Typically, it will be necessary for an organization to invest in deploying information security policies, processes, and technologies in order to protect the confidentiality, C, integrity, I, and availability, A, of its business processes. Defences deployed against each of C, I, and A may compromize the other. For now, we neglect integrity, focussing on trade-offs between confidentiality and availability. This simplification is justifiable: in many — though by no means all — situations, corruption of data is not a major issue, and we can be concerned just with the availability of uncorrupted data. In particular, this assumption is reasonable in the context of the empirical study by Beatement et al. (2008) [2] of the use of USB memory sticks, which is discussed at length above and which provides a partial motivation for the model described herein. Of course, there are many situations in which such an assumption is quite unsustainable: Different instantiations of our modelling framework can, as discussed above, capture such situations.

So, in order to formulate its security policy, an organization must determine its security preferences. That is, for each of its business processes, determine the extent to which it prefers to protect each of C, I, and A. For one example, an online bookstore may prefer to defend the availability of its website in order to

protect revenue. To do so, it may increase the number and geographical distribution of its servers, thereby greatly increasing the attack surface of the system, and so potentially compromising the confidentiality of data held by the system. For another example, a government intelligence service may be prepared to sacrifice system availability in order to protect the confidentiality of its secrets.

In earlier work with other co-authors Beautement et al. (2008) [2], described above, two of us have established some experimental evidence for the existence of a trade-off between availability and confidentiality — integrity was indeed neglected in this context, in which corruption of data is a relatively minor issue — in the use of USB memory sticks by the employees of a large financial services organization.

In the presence of trade-offs between the constituent components of information security, we adopt a well-established analytical methodology employed in macroeconomics to model optimal instrument setting by the monetary authorities (e.g., central banks) when faced with trade-offs between the economic magnitudes that they wish to control, such as inflation and unemployment.

Following Giannoni and Woodford (2002) [3], the general linear stabilization policy problem can be expressed as a solution to the following control problem, in which the economic interaction structure of the state variables is given in terms of a linear system of the form

$$G \begin{bmatrix} Z_{t+1} \\ E_t z_{t+1} \end{bmatrix} = A_1 \begin{bmatrix} Z_t \\ z_t \end{bmatrix} + A_2 r_t + A_3 u_t \tag{1}$$

where z denotes a vector of endogenous variables and the vector of pre-determined variables is given by Z. The instrument available to the authorities is given by r and the system is disturbed from its original equilibrium position due to the existence of shocks u_t. The objective of the policy is to minimize the quadratic objective function in terms of squared deviations of the variables of interest $\mathit{\Pi}$ from some a-priori specified target values $\mathit{\Pi}^*$ by choosing the appropriate value of r given the structure of the system, the loss function,

$$\Lambda = E_t \left\{ \sum_{t=0}^{T} \frac{\delta^{-t}}{2} \left(\mathit{\Pi} - \mathit{\Pi}^* \right)^{\mathsf{T}} \Omega \left(\mathit{\Pi} - \mathit{\Pi}^* \right) \right\} \tag{2}$$

where the vector of variables denoted by $\mathit{\Pi}$ includes values of both z and r. The matrix Ω denotes the variance covariance matrix of the system and δ is the authorities' discount factor. The conditional (on all available information) expectations operator is E_t.

The equilibrium characterization of the system consists of a set of time invariant equations:

$$z_t = \beta_0 + \beta_1 \bar{Z}_t + \beta_2 \bar{u}_t \tag{3}$$

where $\bar{\cdot}$ indicates that the structure of the relevant vectors can differ from the one denoted in Equation 1. The imposition of rational expectations requires that the model's predictions of the endogenous variables are equal to the agents' forecasts.

Nobay and Peel (2003) [8] accommodate the absence of symmetric loss in the presence of deviations by employing, in Λ, the linex function whose asymmetry depends upon the choice of the parameter a:

$$g(x_t) = \{\exp(ax_t) - ax_t - 1\} / a^2.$$

In our case, we restrict our analysis to quadratic loss functions but we allow for unequal weights to be applied to its different arguments.

The analysis given by Giannoni and Woodford (2002) [3], together with refinements of the kind suggested by the work of Nobay and Peel (2003) [8], provides a very general framework for capturing the dynamics of investments and trade-offs in information security within which the choices of security and investment properties to be modelled appropriate for a given context, along with associated organizational preferences, can be captured.

In the next section, we develop a model of this type in the context of information security that is inspired by the study presented in Beautement et al. (2008) [2] and briefly discussed above. For simplicity of analysis, we begin with a continuous time model — a conceptually convenient approximation often employed in many mathematical modelling contexts — which we later discretize. We work with a utility, or loss, function that is quadratic in each of its components.

We then examine the system's response to temporary (one-time) shocks (or perturbations, or disturbances) and map the time evolution of the both the control and state variables. Within this framework we are able to gauge the responses to shocks in terms of magnitude and duration. The stability of the system guarantees the eventual return to a stable path. Such methodology for the examination of the responses of a multivariate linear/non-linear system is well-established in the econometric literature, in the context of linear and non-linear vector autoregressive systems, where the impulse–response function (IRF) is calculated (see Hamilton (1994) [5]). An impulse–response function traces out the response of a state-variable of interest to an exogenous shock (this is normally unobserved). Usually the response is portrayed graphically, with time horizon on the horizontal axis and the magnitude of difference between the undisturbed system and its response to the shock on the vertical axis. Monte Carlo methods are then used for statistical inference to establish whether the calculated responses are statistically significant. In this study, we develop a dynamic system that is subject to a single stochastic disturbance (to confidentiality) and we study the IRF of such system under alternative sets of structural parameters.

3 The Model and Its Meaning

We have explained, in § 2, how we understand confidentiality, integrity, and availability to trade-off against one another. Simplifying, we can neglect integrity — we assume that our storage and processing technologies do not corrupt data — and study the trade-off between confidentiality and availability. This situation is intuitively appealing: disks, DVDs, and memory sticks are quite rarely corrupted, at least in contexts similar to that studied in Beautement et al. (2008)

[2]: increasing a system's availability — for example, by increasing the number and distribution of a system's web-servers — may be thought of as increasing the attack surface of the system, and so reducing the confidentiality of (the information contained within) the system.

The starting point is the utility function, more naturally thought of here as a loss function, expressing the system operators confidentiality and availability preferences. In the given definition, C refers to the aggregate level of confidentiality of information in the system, and $\bar{C}$ is its target, A refers to the aggregate level of availability of information in the system; K denotes the capital stock in information security (i.e., the aggregate value of investments in information security to-date).

We postulate a system whose optimal operational state $(\bar{C}, \bar{A})$ is below its maximal capacity. If the system exceeds such levels the system's reliability becomes problematic and consequently the system's manager attempts to restore it at the predetermined optimal levels. The same happens when the system underperforms because of an 'attack' or any other security breach. The control mechanism in both cases is

$$R = \frac{1}{1-\xi}, \text{for } \xi \in [0,1)$$

which may be thought of as capturing the complexity of the system via the extent to which the system is inter-connected: if the proportion of of the system that is inter-connected is zero (i.e., $\xi = 0$), then the system's complexity is trivial (i.e., 1); as ξ tends to 1, however, the complexity of the system tends to infinity. Such a response aims to alter the system's availability. This may be seen as controlling access to the system.

In addition, we postulate that investment in information security that helps managers to restore the system is expensive, as large deviations from its pre-announced target levels undermine the 'credibility' of the managers and may not be authorized by the CFO. The important element here is the presence of the three elements of deviations form pre-agreed targets in the loss function. Further developments can allow for more sophisticated functional forms that restrict the solutions to one-sided deviations from targets. Notice that, as we measure all metrics in the $(0,1)$ (or $[0,1)$) interval, that is as proportions, the size of the system is assumed constant. This is an area that we may wish to develop in future models by adopting a metric such as capital stock in information security 'per machine' in the network.

The equations below represent the decision-makers' optimal control problem

$$L(C, A, \dot{K}) = E\left(w_1(C-\bar{C})^2 + w_2(A-\bar{A})^2 + w_3(\dot{K}-\bar{\dot{K}})^2\right), \tag{4}$$

the loss function, whose solution will be of the form

$$L(R) \triangleq \min_x L(C, A, \dot{K}) \tag{5}$$

where x is a control variable. In this case, the optimal control issue is based on convex preferences relative to a given set of targets, $\bar{C}$, $\bar{A}$ and $\bar{\dot{K}}$. These are

as follows, the target confidentiality, C, availability, A, and target change in investment in information security, $\dot{K}$.

The weights (w_1, w_2, w_3) represent the type of organization, expressing, as discussed in § 4, the organization's security profile preferences. The time evolution of confidentiality and availability are described in Equations 6 and 7. C_0 is an initial value.

$$C = -\alpha(P)\left(\int_{t_0}^{t} \dot{A}\, dt \;\left(\beta \int_{t_0}^{t'} \dot{K}\, dt'\right)^{-1}\right) + C_0 \tag{6}$$

$$A = \gamma\left(\int_{t_0}^{t'} \dot{R}\, dt'\right) + \delta\left(\int_{t_0}^{t'} \dot{K}\, dt'\right) - \epsilon\left(\int_{t_0}^{t'} \dot{C}\, dt'\right) \tag{7}$$

where $t' < t$.

Investment in information security is triggered by fluctuations in availability and the time dynamics of this are expressed in Equation 8

$$\dot{K} = -\eta \dot{A} \tag{8}$$

The system responds to deviations in confidentiality, as given by Equation 9:

$$\dot{R} = x\left(C - \bar{C}\right) \tag{9}$$

Note that, as $t' \to \infty$, the system stabilizes.

As formulated here, our model shocks only confidentiality. A richer model might, for example, also shock availability. Such a model would need to be formulated with an additional control instrument, so that there would be an instrument corresponding to each shocked dimension.

The weights in the loss function (4) characterize the type of the organization; for example, military and deep-state organizations might put a great deal of weight on C compared to A, whilst a retailer or welfare distributor might place greater value on A compared to C. Finally, the weight on $(\dot{K} - \bar{\dot{K}})^2$ reflects the system's loss when managers are forced to compromize budgets. Public organizations may be more restricted, compared to private sector firms, and therefore be more reluctant to miss $\bar{\dot{K}}$, implying a higher weight associated with this deviation in the loss function. The term $w_3(\dot{K} - \bar{\dot{K}})^2$ deserves more discussion: this is the credibility of the decision maker: if the investment needs to be increased (or decreased) by a large amount, given a conditional set-up, then the initial guesses of the decision maker in setting the equilibrium change in investment were faulty and this results in a subsequent loss of credibility. For example, if a government sets a level of growth in spending of $\bar{\dot{K}}$, then a sudden requirement to increase the level of spending, from time t to $t + \Delta t$ results in ΔK: if $\frac{\Delta K}{\Delta t} \gg \bar{\dot{K}}$, then the decision makers' credibility is decreased (based on convex credibility preferences) with subsequent loss of welfare.

Equation 4 is the objective function, which we seek to minimize, and Equation 5 denotes the solution from the optimization of the control variable x from 9.

See Appendix B for an explanation — in terms of basic utility theory — of the justification of loss functions of this form.

Breaches in confidentiality are denoted by P, the stochastic process that generates such events. Their impact is measured by α, and such breaches will be referred to as shocks to the dynamic system represented in equation 6. The system's attack surface is modelled by the availability

$$\int_{t_0}^{t} \dot{A}\, dt,$$

and amplifies the influence of breaches, whilst increases in the capital stock of information security[4],

$$\frac{1}{\int_{t_0}^{t} \dot{K}\, dt},$$

mitigates against the severity of the shock. The effectiveness of this mitigation is measured by the value of the positive parameter β. The availability of the system depends positively of the system's inter-connectedness, $\int_{t_0}^{t} \dot{R} dt$ and the capital stock of information security. Increases in confidentiality are expected to exert a negative influence on the system's availability. The positive parameters γ, δ, and ϵ measure the impact of these factors on the system's availability.

IT managers will respond to decreases in availability by increasing investment in information security (8). The managers' response is measured by the parameter, η. In the presence of deviations of confidentiality from its target, IT managers respond by manipulating the system's inter-connectedness. Such response is calculated optimally given the architecture of the system, as captured by the parameters α, β, γ, δ, and ϵ, and the managers' preferences and behaviour as captured by w_1, w_2, w_3, and η, given the choice of targets $\bar{C}, \bar{A}$, and $\bar{K}$.

This set-up offers the opportunity to characterize systems according to their architecture and combine them to the preferences of managers. For example, systems with very effective information security capital and managers valuing availability above confidentiality, $w_2 > w_1$, will adjust differently to the same shock in confidentiality if $w_1 < w_2$. In addition to these distinctions differences in behavioural characteristics, η, will determine the relative rate of adjustment. The multi-variate structure of this model, with its general system parameters, is sufficiently expressive to be able to capture a wide range of system profiles of interest, such as the deep-state and commercial systems previously mentioned.

Table 1 illustrates a proposed set of parameter values for three classes of organization: military, financial, and retail. Each organization has varying requirements for its system's robustness to shocks. For example, military-type organizations require confidentiality to be maintained in preference to availability. As such, the α parameter would be expected to be very large, matching the sensitivity of this type of organization to loss of confidentiality. This cost is determined by the ϵ parameter, which is also very high for this class. Also, military

[4] For simplicity of exposition of the initial properties of the model, we do not allow for depreciation in the capital stock of information security.

organizations would be expected to have a high β parameter, given the level of control required relative to the level of expenditure, K. In contrast, financial- and retail-type organizations need to operate on a day to day basis and as such have much higher γ parameters. The cost of loss of reputation to retail is higher and, as such, the main difference between retail organizations and military organizations should be characterized via the ϵ parameter (small for military and very large for retail). Finally, the feedback between change in investment, $\dot{K}$, and change in availability, $\dot{A}$, characterized by the η parameter is also very different for retail and financials, but very similar for military and retail, illustrating financial organizations' ease of redistributing resources for security purposes.

Table 1. Organizational Preferences

Organization Type	System Parameters	Managers' Preference Parameters
Military	$\alpha \gg 0$ $\beta > 0$ $\gamma \to 0$ $\delta < \gamma$ $\epsilon > 0$ $\eta \to 0$	$w_1 \gg w_2 > w_3$
Financial	$\alpha \to 0$ $\beta \to 0$ $\gamma \gg 0$ $\delta \to 0$ $\epsilon \gg 0$ $\eta \gg 0$	$w_1 \simeq w_2 > w_3$
Retail	$\alpha \to 0$ $\beta \to 0$ $\gamma \gg 0$ $\delta \to 0$ $\epsilon \gg 0$ $\eta \to 0$	$w_2 \simeq w_3 \gg w_1$

Having postulated a model — justified by elementary considerations of the nature of investments in information security, including how systems incorporate such investments — we now proceed to examine the system's response to perturbations under alternative parameter constellations that characterize systems and managers with different preferences and behaviours. Table 1 provides examples of the preference of different types of organization.

4 Numerical Examples and Simulations

We select parameter constellations to characterize some systems of interest. We apply the same shock to confidentiality to each of these these systems, and discuss the comparative responses. To proceed with this task, we use the following discretization scheme for the model, with full details given in Appendix A:

$$C_{t+\Delta t} = -\alpha\left(E\left(P_{t+\Delta t}\right)\right)\left(\sum_{t=0}^{t}\Delta A_t + A_0\right)\left(\beta\sum_{t=0}^{t}\Delta K_t + K_0\right)^{-1} + C_0 \quad (10)$$

$$A_{t+\Delta t} = \gamma\left(\sum_{t=0}^{t}\Delta R_t + R_0\right) + \delta\left(\sum_{t=0}^{t}\Delta K_t + K_0\right) - \epsilon\left(\sum_{t=0}^{t}\Delta C_t + C_0\right) \quad (11)$$

$$\Delta K_t = -\eta A_t \quad (12)$$

$$\Delta R_t = x\left(C_t - \bar{C}\right) \quad (13)$$

$$K_{t+\Delta t} = K_t + \Delta K_t \quad (14)$$

$$R_{t+\Delta t} = R_t + \Delta R_t \quad (15)$$

The evolution of the model will be non-explosive provided that the roots of the following polynomial lie within the unit circle:

$$\begin{aligned}\varsigma = {} & Z^5 - Z^4 + \left(-\ln(\epsilon)\ln(\alpha) + \ln(\delta)\ln(\eta)\right)Z^3 \\ & + \left(\ln(\epsilon)\ln(\alpha) - \ln(\theta)\ln(\alpha)\ln(\gamma) + \ln(\epsilon)\ln(\beta)\ln(\eta)\right)Z^2 \\ & + \left(\ln(\theta)\ln(\beta)\ln(\gamma)\ln(\eta) - \ln(\delta)\ln(\eta)\right)Z \\ & + \ln(\theta)\ln(\beta)\ln(\gamma)\ln(\eta) - \ln(\epsilon)\ln(\beta)\ln(\eta)\end{aligned} \quad (16)$$

The full derivation of this stability condition is given Appendix A.

To elucidate the impact of a single non-persistent shock to confidentiality, C_t, the impulse–response of C_t to a shock to P_t at $t = 0$ is derived numerically. For tractability and exposition, the system responses are illustrated as a percentage deviation from equilibrium of the system following a single unit-shock to confidentiality (i.e., we assume that $P_{t=0} = 1$).

We now illustrate the applicability of our model by exploring, in the subsections below, constellations of parameters that characterize contrasting types of organizations (Organization 1, Organization 2). We denote the contrasting choices of parameters by subscripting with 1 and 2: e.g., w_{11}, w_{12}, etc.. It should be noted that, in all cases, the system returns to equilibrium in finite time.

Example 1: Confidentiality versus Availability

$w_{11} \gg w_{12}$, $w_{22} \gg w_{21}$

We compare the behaviour of Organization 1, such as a deep-state or intelligence agency, which weighs confidentiality more highly than availability, with Organization 2, such as an online retailer, which weighs availability more highly than confidentiality. These preferences are expressed by the relative values of w_{11} and w_{12}. We assume, for simplicity, that the organizations are similar in all respects.

Example 2: Impact of Confidentiality Deviations

$\gamma_1 \gg \gamma_2$, $\eta_1 > \eta_2$

We compare two otherwise similar organizations for which the impact of the degree of their network inter-connectedness, and hence of deviations of confidentiality from target, is very different. In Organization 1, the parameter γ_1 is relatively large, so that the impact of deviations of confidentiality from the target is large. Organizations with this characteristic might include banks or health agencies. In contrast, Organization 2, which might be a public information service or a social networking site, γ_2 relatively small, so that deviations of confidentiality below target have a relatively small impact on availability. Since $\eta_1 > \eta_2$, Organization 1's investment response is greater than Organization 2's.

Example 3: Level of Vulnerability and Response

$\alpha_1 < \alpha_2$, $\beta_1 < \beta_2$, $\eta_1 > \eta_2$

We compare to otherwise similar organizations which have different levels of vulnerability. Organization 2 is more vulnerable, as $\alpha_1 < \alpha_2$, which is mitigated by greater investment, $\beta_2 > \beta_1$. However, since $\eta_1 > \eta_2$, Organization 1's investment response is greater than Organization 2's.

The three examples above have been chosen to illustrate the effects of changes in essentially one dimension. Clearly, more realistic comparisons would require more delicate analyses with more variation in the various parameters.

In the three sections that follow below, corresponding to the three examples described above, we plot the impulse–response of the system to a unit shock. In each case, we plot for each comparative pair, the following:

- Deviation from equilibrium of each of C, A, and K;
- The evolution of the control variable, x (recall Equation 9).

Organization 1 is plotted on the left, Organization 2 on the right.

Example 1: Confidentiality versus Availability, Figure 1

The recovery of confidentiality and availability to their pre-shock levels is consistent with the managers' preferences. Measures are taken to restore the system's degree of confidentiality rapidly by enforcing prolonged periods of reduced interconnectedness.

In Organization 1, capital in information security increases almost immediately and then declines monotonically whilst for Orgainzation 2, both confidentiality and availability are restored at almost the same rate whilst capital in information security is of relatively smaller size and it achieves it maximum few periods after shock, exhibiting a somewhat slower rate of return to 'equilibrium'.

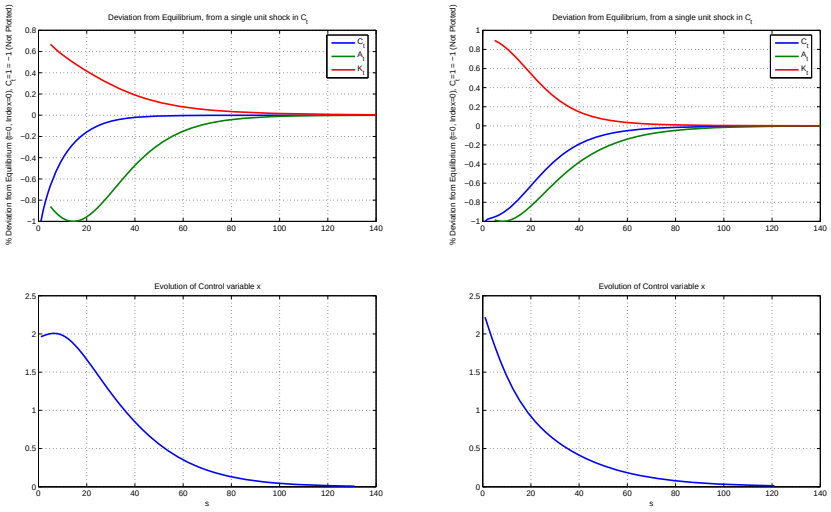

Fig. 1. Confidentiality (w_1) versus Availability (w_2)

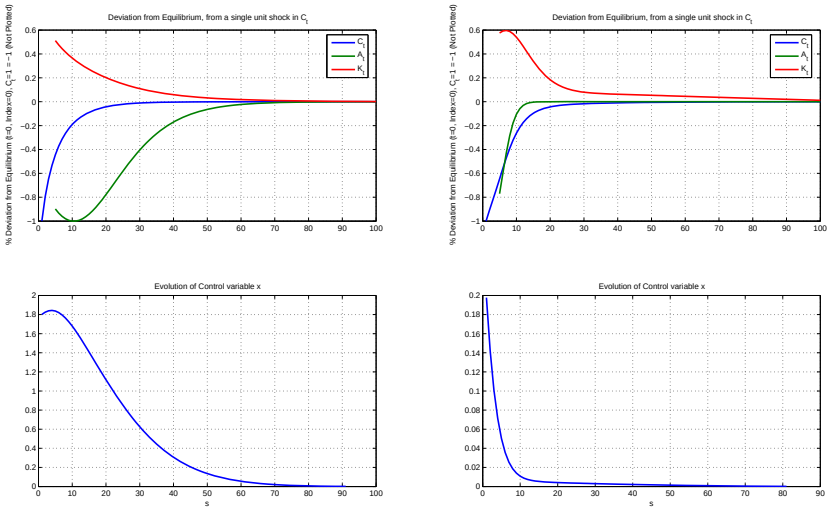

Fig. 2. Impact of Confidentiality Deviations

Example 2: Impact of Confidentiality Deviations, Figure 2

In Organization 1, confidentiality is restored rapidly, and availability lags behind. In Organization 2, confidentiality is restored less rapidly, and availability is the priority. System inter-connectedness is restored less rapidly in the first organization. The evolution of the capital stock is radically different in the two cases. For Organization 1, the initial increase is followed by monotonic reversion to equilibrium, its maximum size not exceeding 0.5. Under the same shock, Organization 2 increases rapidly its capital stock over the subsequent period achieving a maximum value of about 0.75. Having achieved this level, the capital stock is restored to its initial level.

Example 3: Level of Vulnerability and Response, Figure 3

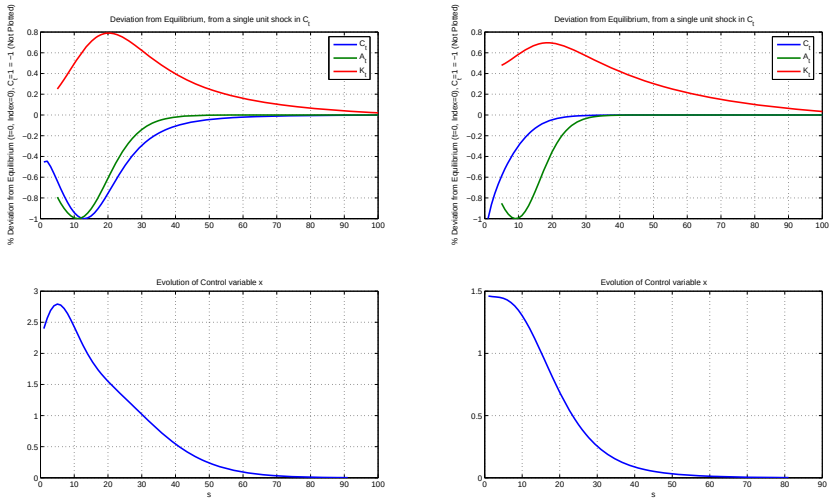

Fig. 3. Level of Vulnerability and Response

Here confidentiality is restored more rapidly in the Organization 2, but availability lags behind. This greater emphasis on security is also reflected by the longer time taken for the second organization to restore system inter-connectedness.

In both cases, the response of capital stock in information security to the shock is not monotonic: they achieve their maxima after approximately 20 periods with Organization 1 exhibiting a modest initial increase followed by subsequent rapid changes bringing the stock of capital well-above the level achieved by Organization 2.

The responses of confidentiality and availability show very different patterns of recovery. The managers' response to the perturbation (the value of x) differs

both in terms of size (in Organization 1 the response to the deviation is far more aggressive) and time evolution, their sensitivity declines fairly rapidly, albeit from a higher base whilst the managers of the second firm maintain for longer periods low levels of system.

5 Conclusions and Directions

We have presented a framework for evaluating the (relative) consequences of C(I)A preferences based on quadratic loss functions.

The following observations, variations, and extensions are suggested:

- A more careful, empirical examination of the assumptions about the systems and management aspects of information security upon which our modelling framework is based;
- More sophisticated forms of loss functions, including asymmetries within and between the confidentiality, availability and investment terms in the loss function — see, for example, the use of linex functions by Nobay and Peel (2003) [8];
- Consideration of the additional dimension of integrity, thus completing the application of our models to the CIA view of information security;
- The model presented here is about a single stochastic threat to confidentiality. Considering multiple threats — with control instruments corresponding to each dimension to which shocks are applied — would strengthen the applicability of the model. Such an extension would require an understanding of the co-variance between threats;
- Different types of investments in information security mitigate against attacks in different ways: for example, we might ditinguish between defences against the likelihood of a breach and defences against the severity of a breach. Such distinctions would, evidently, require refinements to our model;
- Qualitatively different types of threat, such as threats to integrity by data-destroying viruses which might be expected to trigger investments in, for example, patching, would require a significantly more complex model utilizing the ideas discussed above.

Acknowledgements. We are grateful to several of our colleagues, and to Matthew Collinson in particular, for their comments on this work. We are also grateful to the anonymous referees for many comments and observations which have helped us to improve the presentation of this work.

References

1. R. Anderson, R. Böhme, R. Clayton, and T. Moore. Security economics and the internal market. Report to the European Network and Information Security Agency (ENISA), 2007, `http://www.enisa.europa.eu/doc/pdf/report_sec_econ_&_int_mark_20080131.pdf`.

2. A. Beautement, R. Coles, J. Griffin, C. Ioannidis, B. Monahan, D. Pym, A. Sasse, and M. Wonham. Modelling the Human and Technological Costs and Benefits of USB Memory Stick Security. In M. Eric Johnson, editor, *Managing Information Risk and the Economics of Security*. Springer, 2008. Preliminary version available in Proc. WEIS 2008: `http://weis2008.econinfosec.org/papers/Pym.pdf`.
3. M.P. Giannoni and M. Woodford. Optimal Interest-Rate Rules I: General Theory. Working Paper Series 9419, National Bureau of Economic Research, 2002. ISSU 9419, ISSN 0898-2937.
4. L.A. Gordon and M.P. Loeb. The Economics of Information Security Investment. *ACM Transactions on Information and Systems Security*, 5(4):438–457, 2002.
5. J.D. Hamilton. *Time Series Analysis*. Princeton University Press: New Jersey, 1994.
6. K. Hausken. Returns to information security investment: The effect of alternative information security breach functions on optimal investment and sensitivity to vulnerability. *Information Systems Frontiers*, 8(5):338–349, 2006.
7. O. Loistl. The Erroneous Approximation of Expected Utility by Means of Taylor's Series Expansion: Analytic and Computational Results. *American Economic Review*, 66(5):904–910, 1976.
8. R.A. Nobay and D.A. Peel. Optimal Discretionary Monetary Policy in a Model of Asymmetric Bank Preferences. *Economic Journal*, 113(489):657–665, 2003.
9. J. Willemson. On the Gordon & Loeb Model for Information Security Investment. Proc. WEIS 2006: `http://weis2006.econinfosec.org/docs/12.pdf`.

A Discrete time representation and stability of the model

Given the model's system representation,

$$C = -\alpha(P)\left(\int_{t_0}^{t}\dot{A}\,dt\;\left(\beta\int_{t_0}^{t'}\dot{K}\,dt'\right)^{-1}\right)+C_0 \tag{17}$$

$$A = \gamma\left(\int_{t_0}^{t'}\dot{R}\,dt'\right)+\delta\left(\int_{t_0}^{t'}\dot{K}\,dt'\right)-\epsilon\left(\int_{t_0}^{t'}\dot{C}\,dt'\right) \tag{18}$$

$$\dot{K} = -\eta\dot{A} \tag{19}$$

$$\dot{R} = \theta\left(C-\bar{C}\right), \tag{20}$$

assuming simple fixed period time indexing, the discrete time analogues are as follows:

$$C_{t+\Delta t} = -\alpha\left(E\left(P_{t+\Delta t}\right)\right)\left(\sum_{t=0}^{t}\Delta A_t + A_0\right)\left(\beta\sum_{t=0}^{t}\Delta K_t + K_0\right)^{-1}+C_0 \tag{21}$$

$$A_{t+\Delta t} = \gamma\left(\sum_{t=0}^{t}\Delta R_t + R_0\right)+\delta\left(\sum_{t=0}^{t}\Delta K_t + K_0\right) \tag{22}$$

$$-\epsilon\left(\sum_{t=0}^{t}\Delta C_t + C_0\right) \tag{23}$$

$$\Delta K_t = -\eta A_t \tag{24}$$
$$\Delta R_t = x\left(C_t - \bar{C}\right) \tag{25}$$
$$K_{t+\Delta t} = K_t + \Delta K_t \tag{26}$$
$$R_{t+\Delta t} = R_t + \Delta R_t \tag{27}$$

For structural stability,

$$\sum\nolimits_{t=0}^{T>t,T\neq\infty} e' y_t < \infty \tag{28}$$

where e is a unit vector and y_t is the vector evolution of the system equations, C_t, A_t, K_t, and R_t.

Setting the system as a vector problem, and taking logs for linearity and simplifying, the system maybe represented as follows:

$$\begin{aligned}\begin{bmatrix}\log C_{n+1}\\ \log A_{n+1}\\ \log K_{n+1}\\ \log R_{n+1}\end{bmatrix} &= \begin{bmatrix}0 & \log\alpha & -\log\beta & 0\\ \log\varepsilon & 0 & \log\delta & \log\gamma\\ 0 & -\log\eta & 0 & 0\\ \log\theta & 0 & 0 & 1\end{bmatrix}' \begin{bmatrix}\log C_n\\ \log A_n\\ \log K_n\\ \log R_n\end{bmatrix}\\ &+ \begin{bmatrix}0 & 0 & 0\,0\\ 0 & 0 & 0\,0\\ 0 & -\log\eta & 0\,0\\ 0 & 0 & 0\,0\end{bmatrix}\begin{bmatrix}\log C_{n-1}\\ \log A_{n-1}\\ \log K_{n-1}\\ \log R_{n-1}\end{bmatrix}\\ &+ \begin{bmatrix}\log C_0\\ 0\\ -\log\theta + \log C_0\\ 0\end{bmatrix} + \begin{bmatrix}u_n\\ 0\\ 0\\ 0\end{bmatrix}\end{aligned} \tag{29}$$

$$\mathbf{\Pi}_1 = \begin{bmatrix}0 & \log\alpha & \log\beta & 0\\ \log\varepsilon & 0 & \log\delta & \log\gamma\\ 0 & -\log\eta & 0 & 0\\ \log\theta & 0 & 0 & 1\end{bmatrix} \quad \mathbf{\Pi}_2 = \begin{bmatrix}0 & 0 & 0\,0\\ 0 & 0 & 0\,0\\ 0 & -\log\eta & 0\,0\\ 0 & 0 & 0\,0\end{bmatrix} \tag{30}$$

Setting the parameter matrices as a square matrix over the recursion length of the system, the system matrix, $\mathbf{F}$, is

$$\mathbf{F} = \begin{bmatrix}\mathbf{\Pi}_1 & \mathbf{\Pi}_2\\ \mathbf{I} & \mathbf{0}\end{bmatrix} \tag{31}$$

where $\mathbf{I}$ is a 4×4 identity matrix and $\mathbf{0}$ is a 4×4 matrix of zeros.

Taking the matrix polynomial roots of the system matrix $\mathbf{F}$ has the following simplified representation:

$$\begin{bmatrix}0\\ 0\\ 0\\ \varsigma\left(Z\right) = 0\end{bmatrix} \tag{32}$$

where the largest eigenvalue is

$$\begin{aligned}\varsigma = & Z^5 - Z^4 + (-\ln(\epsilon)\ln(\alpha) + \ln(\delta)\ln(\eta))\, Z^3 \\ & + (\ln(\epsilon)\ln(\alpha) - \ln(\theta)\ln(\alpha)\ln(\gamma) + \ln(\epsilon)\ln(\beta)\ln(\eta))\, Z^2 \\ & + (\ln(\theta)\ln(\beta)\ln(\gamma)\ln(\eta) - \ln(\delta)\ln(\eta))\, Z \\ & + \ln(\theta)\ln(\beta)\ln(\gamma)\ln(\eta) - \ln(\epsilon)\ln(\beta)\ln(\eta) \end{aligned} \tag{33}$$

Therefore the stability of the system will be dependent on the roots of the polynomial from (33) being within the unit circle. For simulation purposes, we transform all parameter values by a fixed constant λ to ensure this stability condition is met.

B Concave Utility and Convex Preferences

For a given representation of preferences, $U = f(x)$, with $x \in \mathbb{R}$, for the domain of the function in the interval, $[a, b]$, where $a > b$, then if, for all possible points characterized by the ordering $a < x_1 < x_2 < x_3 < b$, if $f(x_2) \geqslant L(x_2)$, where $L(x)$ is a straight line running through, $(x_1, f(x_1))$ and $(x_3, f(x_3))$, the function is said to be concave in the domain $[a, b]$. This also implies that

$$f'(x_1) > f'(x_2) > f'(x_2) \tag{34}$$

$$f''(x_1) < 0 \tag{35}$$

Consider the second-order Taylor expansion of U,

$$\Im(U) = f(\bar{x}) + \frac{f'(\bar{x})}{1!}(x - \bar{x}) + \frac{f''(\bar{x})}{2!}(x - \bar{x})^2 + r \tag{36}$$

Loistl (1976) [7] determines that for standard maximization problems the remainder term is zero, if x is a random variable $x \in \mathbb{R}$, and the moments of x are uniquely determined by its first non-centralized $E(x)$ and second centralized moment $E(x - \bar{x})^2$. For a general target problem, if we consider the expected value of x to be the target value $\bar{x}$, then the following conditions are assumed in equilibrium,

$$r = 0 \tag{37}$$

$$E(x - \bar{x}) = 0 \tag{38}$$

$$E(x - \bar{x})^2 > 0 \tag{39}$$

$$f''(x_1) < 0 \tag{40}$$

For a given set of control variables Ω, whereby $E(x - \bar{x})^2 | \Omega$, maximum welfare is obtained when

$$\Im(U) \triangleq \max_{\Theta} \Im(U | \Theta) \tag{41}$$

Given that, for all x, $f''(x)$ is negative and monotone decreasing to 0 with increasing x_0, the maximization problem inverts to a loss minimization problem by setting $\frac{1}{2}f''(x) = -w$. Utility maximization occurs when

$$\max_{\Theta} \Im\left(U\left|\Theta\right.\right) \equiv \min_{\Theta}\left(w\left(x-\bar{x}\right)^{2}\left|\Theta\right.\right) \tag{42}$$

B.1 Addition Rules

Consider the variables x, y and z and a representative individual with concave utility $U = f(x,y,z)$, where $(x,y,z) \in \mathbb{R}^3$ for any set of 3-tuple points bounded by $a_{x,y,z} < b_{x,y,z}$; that is,

$$\mathcal{X} = \left\{ \begin{array}{c} a_x < x_1 < x_2 < x_3 < b_x \\ a_y < y_1 < y_2 < y_3 < b_y \\ a_z < z_1 < z_2 < z_3 < b_z \end{array} \right\} \tag{43}$$

The function is concave iff $f(x_2,y_2,z_2) \leq L(x_2,y_2,z_2)$, for all feasible points,

$$\left(f\left(a_x,a_y,a_z\right),a_x,a_y,a_z\right) \tag{44}$$
$$\left(f\left(b_x,b_y,b_z\right),b_x,b_y,b_z\right) \tag{45}$$
$$\left(f\left(x_1,y_1,z_1\right),x_1,y_1,z_1\right) \tag{46}$$
$$\left(f\left(x_2,y_2,z_2\right),x_2,y_2,z_2\right) \tag{47}$$
$$\left(f\left(x_3,y_3,z_3\right),x_3,y_3,z_3\right) \tag{48}$$

where $L(.)$ is the hyperplane that passes through $\left(f\left(x_1,y_1,z_1\right),x_1,y_1,z_1\right)$ and $\left(f\left(x_3,y_3,z_3\right),x_3,y_3,z_3\right)$. Again this implies that each partial second-order derivative of U is negative:

$$\frac{\partial^2 f(x,y,z)}{\partial x^2} < 0 \qquad \frac{\partial^2 f(x,y,z)}{\partial y^2} < 0 \qquad \frac{\partial^2 f(x,y,z)}{\partial y^2} < 0 \tag{49}$$

Again given a vector Taylor expansion around a set of target points $(\bar{x},\bar{y},\bar{z})$,

$$\Im(U) = \sum_{j=0}^{\infty}\left(\frac{1}{j!}\left(\mathbf{a}.\nabla_{\mathbf{r}'}\right)^{j} f\left(\mathbf{r}'\right)\right)_{\mathbf{r}'=\mathbf{r}} \tag{50}$$

and eliminating cross products and setting

$$\frac{\partial^2 f(x,y,z)}{\partial x^2} = -w_x \qquad \frac{\partial^2 f(x,y,z)}{\partial y^2} = -w_y \qquad \frac{\partial^2 f(x,y,z)}{\partial y^2} = -w_z \tag{51}$$

and given x, y, and z are independent randomly distributed random variables, uniquely defined by their first two moments, the utility maximization problem inverts to the following loss minimization function

$$\max_{\Theta}\left(\Im(U)\right) \equiv \min_{\Theta}\left(w_x\left(x-\bar{x}\right)^2 + w_y\left(y-\bar{y}\right)^2 + w_z\left(z-\bar{z}\right)^2\right) \tag{52}$$