Fixed Costs, Investment Rigidities, and Risk Aversion in Information Security: A Utility-theoretic Approach

Christos Ioannidis* David Pym† Julian Williams‡

Abstract

This paper addresses the question of determining the optimal timing of interventions in information security management. Using utility theory, we derive the limiting condition under which, given a potential or realized risk, a decision to invest, delay, or abandon can be justified. Our primary focus is on the decision to defer costly deterministic investments, such as the removal of a service or implementation of a security patch, when the costs associated with future security vulnerabilities are uncertain. We outline an investment function with irreversible fixed costs that introduces a rigidity into the investment decision-making profile. This rigidity introduces delay in the implementation of security measures, resulting in cyclical investments in information security, as the decision-maker determines the optimal investment horizon. We therefore show that cycles emerge endogenously given the policy-maker's chosen trade-offs between investment and the deterioration of the system attributes.

1 Introduction

Decision-making in information security problems has traditionally been understood in the context of the equalization of marginal cost of, and marginal benefit from, investments. An alternative approach is to formulate the same problems in terms of strategic behaviour, such as Bayesian games. For a classic example among many, consider Fultz and Grossklags' paper on distributed security attacks [15]. This approach is useful for garnering certain stylized facts about behaviour that can guide the formulation of security policies and the design of incentive structures for policy-makers in information systems. Because these models rely upon high levels of abstraction from the detail of the system's design and its security management, such approaches are not easily mapped to deployable decision-support tools. In this paper, we argue that the appropriate formulation of the decision-support problem is as a decision-under-uncertainty problem for a representative policy-maker.

We formulate this optimization as a utility maximization problem under uncertainty; that is, framing the decision problem as a one-dimensional single-equation reaction function, with multivariate stochastic innovations in the system architecture, representing the risks that undermine the system attributes. This is, in effect, an Arrow–Pratt [6, 25] approach to optimal decision making under uncertainty for a risk-averse policy-maker. Using this approach, we are able to derive a closed-form limiting condition for decision-making using all the information and intervention technologies that are available to the decision-maker.[1]

The assumption of risk aversion is justified as a reasonable choice because: in the context of government, a precautionary approach to risk management is the usual guiding principle; in owner-manager businesses, the capital of the firm is intimately connected to the capital of the

*Department of Economics, University of Bath, Bath BA2 7AY, England, UK; c.ioannidis@bath.ac.uk.

†School of Natural and Computing Sciences, King's College, University of Aberdeen, Aberdeen AB24 3UE, Scotland, UK; d.j.pym@abdn.ac.uk. Part of this work was carried out whilst Pym was employed at HP Labs, Bristol.

‡Business School, University of Aberdeen, Aberdeen AB24 3QY, Scotland, UK; julian.williams@abdn.ac.uk.

[1]The solution includes the stochastic process that represents the threat environment.

policy-maker; in businesses in which management and ownership are separated, transaction costs associated with change of employment inhibit the policy-maker's risk appetite.

An economics-based perspective on the management of information security operations has been discussed in a series of papers [2, 4, 3]. In [17, 18], a microeconomic analysis of the costs and benefits of defences against given vulnerabilities is presented. Recent work by the present authors has applied ideas from utility theory and dynamic optimization to information security. More specifically, we have presented a dynamic model of trade-offs between confidentiality, availability, and investment in information security operations. [20, 19].[2] The aggregate timing of decision-making — in product cycles, consumption cycles, and investment cycles — is explored in, for example, the recent work of Bloom [10].

In the case of risk management and mitigation of vulnerability, this is a utility-of-action problem, which has been explored extensively in applied, commercial, contexts in [8, 9, 24], work which has directly informed the present paper. To demonstrate this issue appropriately in a timing framework, we treat system vulnerabilities as being stochastic discount factors that erode the attributes of a system. The discount factors are assumed to be arrive through a poisson process and have jump intensities driven by a multivariate log-normal distribution. Whilst this is a simple specification, the purpose of this paper is to demonstrate a methodology and the steps involved in implementing it.

In this paper, we introduce the notion of a security investment cycle. The key components of this idea are the following:

- Investments in IT systems are a major cost for firms;
- These systems are subject to security vulnerabilities, which typically compound over time;
- Accordingly, IT system managers plan investments in information security operations;
- These investments typically include irreversible (e.g., up-front) fixed costs;
- This rigidity (e.g., derived from up-front costs) inhibits rapid implementation of new security measures in response to emerging threats, since the fixed costs may outweigh the immediate benefits;
- To exhibit this situation, consider how the two system attributes of performance and security trade off against each other, and jointly against investment costs.

An example of this situation is provided by the question of when to deploy patches. Indeed, the importance of timing vulnerability management in networks in the presence of externalities has been addressed by [7]. They show that software vendors can offer rewards to encourage timely patching when vulnerabilities occur in both proprietary software and freeware and, given the differential costs of patching to users, conclude ([7], p. 1718) that 'a "one-size-fits-all" approach is unlikely to be an immediate remedy'. In addition, for situations in which the actions of users impact upon the welfare of others, they develop a set of incentive structures for the implementation of effective patch management. The timing of vulnerability disclosures by vendors is modelled by [5], where it is shown that, with no regulation, the vendor releases a patch less frequently than is socially optimal. In [11], the relationship between the release of patches by vendors and their implementation by users is studied. They classify patching cycles into time-driven and event-driven. They show that social loss is minimized when vendor releases are synchronized with the time-driven cycles of the system operator. When such synchronization cannot be achieved because it is costly, the imposition of liability on the vendor cannot achieve the socially timing optimal of disclosures. Finally, in [11], the authors calculate the socially optimal window of exposure and decompose the patching process into time- and event-driven incidents in a game-theoretic setting. In [16], a financial model using real options — in particular, deferment options — is employed to offer an integrated framework encompassing the decision to delay patching in the presence of known vulnerabilities. When system operators employ a variety of applications, patch arrivals

[2]Modelling multiple trade-offs can be accommodated within the same methodology.

will appear as random events, without apparent periodicity. From previous studies described, it is apparent that the time of patch deployment is important because such an action involves costs and mistiming exacerbates their impact.

The methodology suggested in this paper can be described succinctly in the following steps: first, we map the preferences of decision-makers to a well specified utility function defined in terms of the system's attributes and control variables[3], such as confidentiality, integrity, and availability, or sensitivity and criticality; second, we derive the moment expansion of the utility function, using a Taylor expansion; third, we postulate the stochastic risk environment and substitute the moments of this process into the higher-order terms of the system's attributes, in addition to the moments of the control variable; finally, we compute the maximum of the expected utility by adjusting the properties of the control variable.

In this paper, we consider the security investment cycle more generally. In Section 2, we explain how some fundamental utility theory can be used to formulate the preferences of a representative policy-maker choosing forward-looking investment profiles. In Section 3, we derive an approximation to the solution space for expected utility maximization for a given functional form of investment in the presence of stochastic threats to system security. In Section 4, we set out a specific stochastic threat space and solve for the equilibrium timing of investment. We illustrate three cases of investment implementation: address vulnerabilities on arrival; never invest; and delay investment for a finite time.

The remainder of the paper employs a range of technical terms defined, for our purposes, in Table 1.

Table 1: Glossary of Terms

Phrase	**Definitions**
Vulnerability	A flaw in the architecture of an information system that may lead to a compromise of, for example, CIA or CS (criticality, sensitivity)
Exploit	A malicious tool that facilitates the exploitation of a vulnerability
Mitigation	A system update that mitigates the vulnerability, partially or completely impairing the function of an exploit

2 A Utility Theory of Vulnerability Management

Our first goal is to orientate the vulnerability management problem in an expected utility-maximization framework. We seek to construct a objective function, whose solution at the maximum is equivalent to the expected utility maximization condition. We state the policy-maker's objective function as

$$\mathbb{E}\left(\mathfrak{U}\left(t,T\right)\right) \triangleq \max_{K(t)} \int_{t}^{T} e^{-\beta t} u\left(x(t);K(t)\right) d\mathbb{P}\left(\omega\left(t\right)\right) \tag{1}$$

where

- T is the terminal time,
- $K(t)$ is a choice of investment function,
- $x(t) = \{x_1, \ldots, x_n\}$ is a n-vector of real-valued system attributes that is stochastic, because of threats, defined over the probability space $(\Omega, \mathcal{F}, \mathbb{P})$ [26],

[3]Well specified, in this case, implies that preferences are represented by this utility function are consistent with rational choice, for a full exposition of axiomatic utility theory and decision making see [14].

- $u(x(t); K(t))$ is an instantaneous real-valued twice-differentiable utility function over the system attributes x, with exogenous parameters the investment function, $K(t)$,
- β is a global discount rate, and
- $\omega(t) \in \Omega$ is an experiment in the probability space $(\Omega, \mathcal{F}, \mathbb{P})$ [26].

Here the idea is that we vary the investment function $K(t)$ in order to maximize expected utility at time t by choosing a future investment time $t^* \geq t$.

Equation 1 provides a general characterization of a variety investment problems. As such, it is difficult to derive general analytic solutions and so we reduce the problem space to a polynomial approximation of Equation 1 for which solutions can be found.

In this paper, we assume a risk-averse policy-maker. In the case of a risk-neutral policy-maker, our analysis collapses to a polynomial approximation to the real options solution for the investment timing problem [28].

2.1 The Power Utility Family

We explore the general problem described above in the case in which $n = 2$. This is the simplest case that fully illustrates our approach. Examples of this case would include the security attributes confidentiality and availability, and we have explored, in less generality, the way in which these attributes trade off against each other elsewhere [20, 19]. In [20], for example, we exogenously imposed an investment cycle on the representative firm within the model. In contrast, in this paper, we demonstrate how an investment cycle arises from investment rigidities.

In economics and finance, the power utility family of functions is the predominant mechanism for defining preferences for inter-temporal decision problems. Whilst for most of our derivation we are agnostic to choice of utility function (our interest is restricted to the ratio of the derivatives), some discussion of the higher level of functional form is relevant and useful for future applied work. The basic power utility construct for a consumption variable $x_i \in \{x_1, x_2\}$, suppressing the control variable K, has a partial utility function defined as

$$u_i\left(x_i\right) = \frac{x_i^{1-\gamma_i}}{1-\gamma_i} \tag{2}$$

where γ_i is the coefficient of relative risk aversion $\mathfrak{R}$, for the i^{th} attribute. Combining the partial utility functions with cross power utility would yield and over utility function of

$$u\left(x_1, x_2\right) = \frac{1}{1-\gamma_1} x_1^{1-\gamma_1} + \frac{1}{1-\gamma_2} x_2^{1-\gamma_2} + 2\frac{\left|x_1 x_2\right|^{1-\gamma_{12}}}{1-\gamma_{12}} \tag{3}$$

Several extensions of the power utility have been proposed in the literature and several of these innovations have useful interpretations for information security problems. From this point onward, for ease of exposition, we shall concentrate on the partial utility functions. Kahneman and Tversky [21] suggest the inclusion of a fixed point, or kink point, k, to discriminate between aversion to risk of loss and aversion to risk of gain. The power utility representation of this approach is

$$u_i\left(x_i\right) = \begin{cases} \frac{1}{1-\gamma_i} x_i^{1-\gamma_i} & \forall x_i > k \\ \\ \frac{1}{1-\tilde{\gamma}_i} x_i^{1-\tilde{\gamma}_i} & \forall x_i \leq k \end{cases} \tag{4}$$

where $\tilde{\gamma}_i \neq \gamma_i$. The inclusion of the fixed point adds a significant complication to the type of optimization suggested herein as the derivatives of $u(x)$ are now discontinuous.

An alternative augmentation is to include a utility profile of the consumption of system attributes at some future point in time. This nesting of future utility allows for a substitution between current expected utility and future expected utility and has been used extensively since first being suggested in [13]. The power utility form is compactly presented as

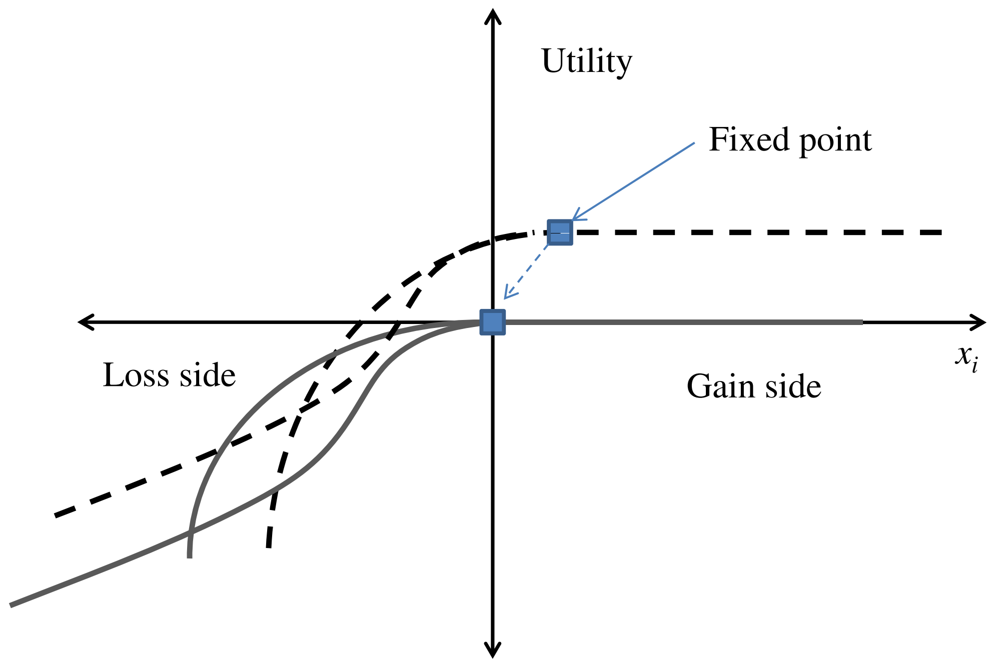

Figure 1: Illustration of the projection of a family of utility functions, $u(x_i)$, for a single attribute. The markers represent the fixed points in the utility problem. The fixed points can be located anywhere within this plane. For example, the dashed represents a curve with a fixed point at at positive value of x_i. For our purposes, we assume the fixed point is at the origin (the dark grey line) and that deviations from steady state are always to the left of the origin.

$$u_i\left(x_i(t)\right) = \left(1-\zeta_i\right) x_i^{\frac{1-\gamma_i}{\theta_i}} + \zeta_i \mathbb{E}_t\left(u\left(x_i(t+\Delta t)\right)^{\frac{1}{\theta_i}}\right)^{\frac{\theta_i}{1-\gamma_i}} \tag{5}$$

where θ_i is the anticipated future coefficient of relative risk aversion at $t+\Delta t$, ζ_i is the intertemporal elasticity of substitution — that is, the substitution between current and future expected utility.

The last type utility function we have considered in our applied work is the 'inside and outside of habit' utility function, suggested by [1]. This sets expected utility as being relative to a peer group represented by an index (of consumption), ξ_i, of the variable x_i. In our notational scheme, the power utility version of this type of utility function is (as usual, suppressing $K(t)$) defined as

$$u_i\left(x_i\right) = \frac{\left(x_i \xi_i^{-1}\right)^{1-\gamma}}{1-\gamma} \tag{6}$$

There are obvious circumstances where each of these definitions of preferences will be appropriate. Augmentations to cater for non-zero cross products (i.e., supermodularity or submodularity) are also relatively trivial. For instance, fixed points are common in many aspects of information security: in particular, on the loss side — essentially, improvements over targets are relatively under rewarded.

Figure 1 outlines an example of this structure.

2.2 The Policy-maker's Problem

We begin with a quick review of several key results for properties of utility functions and in particular the structure of risk aversion. Absolute risk aversion ($\mathfrak{A}$) in two-variable (i.e., $n = 2$) multi-attribute decision problems is defined as follows:

$$\mathfrak{A}(x_1, \cdot) = \frac{u''_{x_1}(x_1, \cdot)}{u'_{x_1}(x_1, \cdot)} \qquad \mathfrak{A}(\cdot, x_2) = \frac{u''_{x_2}(\cdot, x_2)}{u'_{x_2}(\cdot, x_2)} \tag{7}$$

where we we suppress in the notation the exogenous parameter, $K(t)$. This is then simply mapped to a relative risk aversion ($\mathfrak{R}$) context

$$\mathfrak{R}(x_1, \cdot) = \frac{-\gamma_{x_1} u''_{x_1}(x_1, \cdot)}{u'_{x_1}(x_1, \cdot)} \qquad \mathfrak{R}(\cdot, x_2) = \frac{-\gamma_{x_2} u''_{x_2}(\cdot, x_2)}{u'_{x_2}(\cdot, x_2)} \tag{8}$$

where γ_{x_1} and γ_{x_2} are the coefficients of relative risk-aversion (i.e., the marginal rate of change in risk-aversion with respect to $u(\cdot)$) for each of the system attributes.

Both $\mathfrak{A}$ and $\mathfrak{R}$ are useful tools in summarizing the properties of specific utility functions: in addition to the risk aversion properties, the cross products for the attributes are useful in elucidating the preference structure. In the bivariate context, there are three main combinations. Consider the following decomposition

$$u(x_1, x_2) = u_1(x_1) + u_2(x_2) + u_{12}(x_1, x_2) \tag{9}$$

where $u(x_1)$ and $u(x_2)$ are the partial utility functions with respect to system attributes x_1 and x_2 and $u_{12}(x_1, x_2)$ is the joint utility adjustment. In the general form of our modelling framework we maintain the general assumption that $u_{12}(x_1, x_2) \neq 0, \forall \{x_1, x_2\} \in \mathbb{R}^2$. For our final analytic solutions, however, we have assumed *separable additivity*; that is,

$$u''_{x_1,x_2}(x_1, \cdot) = 0 \quad \forall x_2 \qquad u''_{x_1,x_2}(\cdot, x_2) = 0 \quad \forall x_1 \tag{10}$$

The contrasting assumptions that maybe made on the shape of the multi-attribute utility function are *supermodularity* whereby

$$u''_{x_1,x_2}(x_1, \cdot) > 0 \quad \forall x_2 \qquad u''_{x_1,x_2}(\cdot, x_2) > 0 \quad \forall x_1 \tag{11}$$

and *submodularity* whereby

$$u''_{x_1,x_2}(x_1, \cdot) < 0 \quad \forall x_2 \qquad u''_{x_1,x_2}(\cdot, x_2) < 0 \quad \forall x_1 \tag{12}$$

Discussion of the appropriate application of these properties is usually driven by game-theoretic models of incentives. For instance, most problems can be treated as separably additive, and as such the attributes rolled in a single linear function. However, in the authors' experience of working with industry and government, compound attacks on multiple system attributes are often more damaging than attacks (of similar component-wise magnitude) that occur at different times. In this case, utility functions incorporating a degree of supermodularity would be most appropriate for describing policy-maker preferences. Cases of submodular preferences are much rarer, although not unheard of. For instance, in a confidentiality, integrity, and availability (CIA) framework, a distributed denial of service (DDOS) attack mixed with a breach of confidentiality could, for certain institutions such as retailers, be understood as being submodular: to some extent, the DDOS mitigates the effectiveness of the confidentiality attack as the system's availability (to the confidentiality attacker) is compromised. A fuller exploration of supermodular and submodular preferences is deferred to future work.

3 Constructing the Expected Utility Function

For simplicity of exposition, we now simplify the decision under uncertainty problem to a policy maker choosing a forward looking investment profile from an initial time t_0; that is, at a point where no existing vulnerabilities are present. The resulting expected timing of investment $t^* > t_0$ is the ex-ante expected amplitude of the investment cycle. Future work will address the 'steady-state' equilibrium investment horizon at time $t > t_0$.

For a given choice of utility function $u : \mathbb{R}^n \to \mathbb{R}$ operating over $n = 2$ system attributes — consumption variables in an economic context — the dynamic representation of the utility function is defined from the terms of the Taylor expansion as

$$\begin{aligned} u\left(x_1(t_0)+\Delta x_1, x_2(t_0)+\Delta x_2\right) = \; & u\left(x_1(t_0), x_2(t_0)\right) + \\ & \left(u_{x_1}\left(x_1(t_0), x_2(t_0)\right)\Delta x_1 + u_{x_2}\left(x_1(t_0), x_2(t_0)\right)\Delta x_2\right) + \\ & \tfrac{1}{2}u_{x_1,x_1}\left(x_1(t_0), x_2(t_0)\right)\left(\Delta x_1\right)^2 + \\ & \Delta x_1 \Delta x_2 \left(u_{x_1,x_2}\left(x_1(t_0), x_2(t_0)\right)\right) + \\ & u_{x_2,x_2}\left(x_1(t_0), x_2(t_0)\right)\left(\Delta x_2\right)^2 \end{aligned} \tag{13}$$

where $x_1(t_0)$ and $x_2(t_0)$ denote initial values, which is a valid approximation as Loistl [23] demonstrates that under fairly mild conditions the remainder converges to zero.

Assuming that the moment generating process is fully described by its first two moments, the following notation applies:

$$\mu_{x_1}(t) = \mathbb{E}_t\left(x_1(t) - \bar{x_1}\right) \tag{14}$$

$$\mu_{x_2}(t) = \mathbb{E}_t\left(x_2(t) - \bar{x_2}\right) \tag{15}$$

$$\sigma_{x_1}(t) = \mathbb{E}_t\left(x_1(t) - \bar{x_1}\right)^2 \tag{16}$$

$$\sigma_{x_2}(t) = \mathbb{E}_t\left(x_2(t) - \bar{x_2}\right)^2 \tag{17}$$

$$\sigma_{x_1,x_2}(t) = \mathbb{E}_t\left(x_2(t) - \bar{x_2}\right)\left(x_1(t) - \bar{x_1}\right) \tag{18}$$

where $\bar{x_1}$ and $\bar{x_2}$ are long-run targets and $\mathbb{E}_t$ is the instantaneous expectation at time t. Substituting these into the utility function above results in the following expected utility function:

$$\begin{aligned} \mathbb{E}\left(u\left(x_1(t), x_2(t)\right)\right) = \; & u\left(x_1(t_0), x_2(t_0)\right) + \left(u_{x_1}\left(x_1(t_0), x_2(t_0)\right)\mu_{x_1} + u_{x_2}\left(x_1(t_0), x_2(t_0)\right)\mu_{x_2}\right) + \\ & \tfrac{1}{2}u_{x_1,x_1}\left(x_1(t_0), x_2(t_0)\right)\sigma_{x_1}(t) + \sigma_{x_1,x_2}(t)\left(u_{x_1,x_2}\left(x_1(t_0), x_2(t_0)\right)\right) + \\ & u_{x_2,x_2}\left(x_1(t_0), x_2(t_0)\right)\sigma_{x_2}(t) \end{aligned} \tag{19}$$

Assuming the existence of threats that degrade the system, induce utility losses, and continuously compound, and which are such that, for all t, $x_1(t) \geq 0$ and $x_2(t) \geq 0$, then the utility function will obey

$$u\left(\bar{x_1}, \cdot\right) \geq u\left(\bar{x_1} + x_1(t), \cdot\right) \qquad \forall t \tag{20}$$

$$u\left(\cdot, \bar{x_2}\right) \geq u\left(\cdot, \bar{x_2} + x_2(t)\right) \qquad \forall t \tag{21}$$

where $\cdot$ is a placeholder in the function. This results in decreasing marginal utility with respect to loss:

$$\frac{\partial u\left(\bar{x_1}, \cdot\right)}{\partial x_1} \geq \frac{\partial u\left(\bar{x_1} + x_1(t), \cdot\right)}{\partial x_1(t)} \qquad \forall t \tag{22}$$

$$\frac{\partial u\left(\cdot, \bar{x_2}\right)}{\partial x_2} \geq \frac{\partial u\left(\cdot, \bar{x_2} + x_2(t)\right)}{\partial x_2} \qquad \forall t \tag{23}$$

We define the following policy parameters, as described in Table 2:

$$w_{x_1} = -u_{x_1}(x_1(t_0), x_2(t_0)) \quad (24)$$
$$w_{x_2} = -u_{x_2}(x_1(t_0), x_2(t_0)) \quad (25)$$
$$v_{x_1} = -2u_{x_1,x_1}(x_1(t_0), x_2(t_0)) \quad (26)$$
$$v_{x_1} = -2u_{x_2,x_2}(x_1(t_0), x_2(t_0)) \quad (27)$$
$$v_{x_1,x_2} = -u_{x_1,x_2}(x_1(t_0), x_2(t_0)) \quad (28)$$

Each of these has a simple interpretation, as described in Table 2.

Table 2: Policy Parameters

Parameter	**Description**
w_{x_1}	Policy weighting applied to first system attribute
w_{x_2}	Policy weighting applied to second system attribute
v_{x_1}	Sensitivity (risk aversion) to variance in first system attribute
v_{x_2}	Sensitivity (risk aversion) to variance in second system attribute
v_{x_1,x_2}	Sensitivity to covariance first and second system attributes

From the asymmetric preference structure, the policy-maker's problem can be expressed as maximizing an expected utility function. The expected utility from the state of the system attributes is defined by the following integral that represents the cost of inaction:

$$\mathfrak{U}(t_0, T|w_{x_1}, w_{x_2}, v_{x_1}, v_{x_2}, v_{x_1,x_2}) = \int_{t_0}^{T} e^{-\beta t}\ell(t|w_{x_1}, w_{x_2}, v_{x_1}, v_{x_2}, v_{x_1,x_2})\,dt \quad (29)$$

$$= \int_{t_0}^{T} e^{-\beta t}\left(w_{x_1}\mu_{x_1}(t) + w_{x_2}\mu_{x_2}(t) + v_{x_1}\sigma_{x_1}(t) + 2v_{x_1,x_2}\sigma_{x_1,x_2}(t) + v_{x_2}\sigma_{x_2}(t)\right)dt$$

where

$$\ell(t|w_{x_1}, w_{x_2}, v_{x_1}, v_{x_2}, v_{x_1,x_2}) = w_{x_1}\mu_{x_1}(t) + w_{x_2}\mu_{x_2}(t) + v_{x_1}\sigma_{x_1}(t) + 2v_{x_1,x_2}\sigma_{x_1,x_2}(t) + v_{x_2}\sigma_{x_2}(t) \quad (30)$$

The additional separable component in the policy-maker's loss function is defined with respect to the additional investment required in the presence of disclosed vulnerabilities.

The objective is to find the policy-maker's cycle time to investment; that is, the upper limit of integration, T, which satisfies the equality of utility of action and utility of inaction. We denote this value as t^*.

3.1 The Investment Function

We assume the following investment function:

$$K(t) = \begin{cases} 0 & \text{if } t < t^* \\ \hat{K} + K_0 e^{\delta t} & \text{if } t = t^* \end{cases} \quad (31)$$

where $\hat{K}$ is the fixed cost, K_0 is the initial investment, and δ is the growth rate in variable cost, and, as defined above, t^* is the timing of the future investment. We assume that the investment function has a deterministic schedule: that is, there is no stochastic variation in the investment profile.

The policy-maker's credibility is defined in terms of deviation from a preset investment profile. That is, the policy-maker's utility is eroded by deviations from target investment. Specifically, we assume the following utility function for K:

$$u\left(\bar{K}\right) \geq u\left(\bar{K}+K\left(t\right)\right) \qquad \forall t \tag{32}$$

$$\frac{\partial u\left(\bar{K}\right)}{\partial K} \geq \frac{\partial u\left(\bar{K}+K\left(t\right)\right)}{\partial K} \qquad \forall t \tag{33}$$

Note that the utility of investment is always lower for deviations from the target investment profile, and that the change in utility is increasing with $K(t)$.

Taking the series expansion and integrating yields the cost of action:

$$\mathfrak{U}_K\left(t_0, T|w_K, v_K\right) = \int_{t_0}^{T} e^{-\beta t}\left(w_K \mu_K\left(t\right) + v_K \sigma_K\left(t\right)\right) dt \tag{34}$$

where, for initial investment K_0, $w_K = -u_K\left(K_0\right)$, $v_K = -2u_{KK}\left(K_0\right)$, $\mu_k(t) = \mathbb{E}(K(t) - \bar{K})$ and $\sigma_k(t) = \mathbb{E}(K(t) - \bar{K})^2$. The complete policy-maker's problem brings together the two additively separable components of the loss function:

$$\mathfrak{D}\left(t_0, T|w_{x_1}, w_{x_2}, v_{x_1}, v_{x_2}, v_{x_1,x_2}\right) = \mathfrak{U}\left(t_0, T|w_{x_1}, w_{x_2}, v_{x_1}, v_{x_2}, v_{x_1,x_2}\right) - \mathfrak{U}_K\left(t_0, T|w_K, v_K\right) \tag{35}$$

In the case of our system attributes, preferences, and investment function, we have

$$\begin{aligned}\mathfrak{D}\left(t_0, T \mid \cdot\right) \quad = \quad & \int_{t_0}^{T} e^{-\beta t}\left(w_{x_1}\mu_{x_1}\left(t\right) + w_{x_2}\mu_{x_2}\left(t\right) + v_{x_1}\sigma_{x_1}\left(t\right) + 2v_{x_1,x_2}\sigma_{x_1,x_2}\left(t\right) + v_{x_2}\sigma_{x_2}\left(t\right)\right) dt \\ & -\int_{t_0}^{T} e^{-\beta t}\left(w_K\mu_K\left(t\right) + v_K\sigma_K\left(t\right)\right) dt\end{aligned} \tag{36}$$

The components of the decision structure in Equation 35 have the following interpretation: the first denotes the cumulative loss from undertaking no additional investment, and the second is the cumulative loss incurred because of the extra expenditure, in the presence of disclosed vulnerabilities.

Equating utility of action and utility of inaction in Equation 35 yields:

$$\mathfrak{D}\left(t_0, T|w_{x_1}, w_{x_2}, v_{x_1}, v_{x_2}, v_{x_1,x_2}\right) = 0 \tag{37}$$

Notice that the utility function given in Equation 1 uses a Lebesgue integral to handle the stochastic structure of the problem, whilst in Equation 29 the same problem is expressed, via the second-order approximation of the utility function, as a finite integral in terms the first two moments of the underlying distribution.

The decision system represented in Equation 35 is a general device for a wide range of optimal control problems depending on the choice of moment functions, $\mu_K(t)$, $\mu_B(t)$, $\mu_{x_2}(t)$, $\sigma_K(t)$, $\sigma_{x_1}(t)$, $\sigma_{x_2}(t)$ and preferences w_K, w_{x_1}, w_{x_2}, v_K, v_{x_1}, v_{x_2}, v_{x_1,x_2}.

4 The Threat Environment and Decision-making

We have developed the decision-making rule without taking into account explicitly the stochastic nature of vulnerability disclosure. We now proceed to give a detailed description of these stochastic processes, and integrate them into the policy-maker's decision-making algorithm.

The model is explicitly derived in terms of pre-defined target levels, $\bar{x_1}$ and $\bar{x_2}$. The impact of disclosure of vulnerabilities degrades these attributes by a factor Y, which is driven by an underlying vector stochastic process $y(t)$. The key feature of this setting is that the longer a system is left vulnerable, the greater the degradation to the system attributes. The stochastic process, $y(t)$, that characterizes the vulnerabilities/threats, is described by the following stochastic differential equation:

$$dy(t) = (\mu(y(t))\,dt + \sigma(y(t))\,dW_t)\,dJ_t \tag{38}$$

where $\mu(\cdot)$ is a vector/matrix function of the driving factors, $\sigma(\cdot)$ is a càdlàg variance generating function and, as such, discontinuities are entirely driven by the one-dimensional jump process J_t, with time-invariant intensity parameter, or rate, λ. This form provides for a whole family of jump processes with random step sizes. We can think of this as being a compound of a smooth variational system (in the Brownian motion) and a discontinuous jump model driven by a poisson process. We now define μ_1 and μ_2 as being the time-homogenous drift-terms, and σ_1, σ_2 and σ_{12} are the time-homogenous elements of the covariance matrix of the underlying multivariate normal random variable driving the jump sizes of $y(t)$; see Appendix A.

The process is driving a continuous stochastic discount factor, a very useful way of approaching the timing problem, as in essence we are comparing the evolution of a stochastic discount rate (for the potential and realized vulnerabilities) and a deterministic discount rate for the cost function.

The evolution of the individual system attribute x_i is written in integral form as

$$x_i(t_0, T) = x_i(t_0)\exp\left(\int_{t_0}^{T} y_i(t)\,dt\right) \tag{39}$$

The terminal distribution (at time $t + \Delta t$) is assumed to be log-normal with compounded jumps (see [27]). The useful properties of this terminal distribution are that the moments can be defined explicitly as products of the moments of the log-normal distribution and the intensity parameter, λ, from the jump process, as given in Equation 38.

We now explain the timing solution space for the expected loss equation. First we find the derive the form of the solution (Proposition 1), and then we establish the existence of solutions (Theorem 1).

Proposition 1: Form of Equilibrium Solution

Let t^* be the expected time of investment. The equilibrium decision function (at t_0) in expectations is the following:

$$\begin{aligned}
\mathfrak{D}(t_0, T|w_{x_1}, w_{x_2}, w_k, v_{x_1}, v_{x_2}) &= e^{-t\beta}x_2(t)\left(\frac{e^{e^{2\mu_2+\sigma_2^2}\left(-1+e^{\sigma_2^2}\right)t\lambda}v_{x_2}}{-\beta+e^{2\mu_2+\sigma_2^2}\left(-1+e^{\sigma_2^2}\right)\lambda} + \frac{e^{e^{\mu_2+\frac{\sigma_2^2}{2}}t\lambda}w_{x_2}}{-\beta+e^{\mu_2+\frac{\sigma_2^2}{2}}\lambda}\right) \\
&\quad + e^{-t\beta}x_1(t)\left(\frac{e^{e^{2\mu_1+\sigma_1^2}\left(-1+e^{\sigma_1^2}\right)t\lambda}v_{x_1}}{-\beta+e^{2\mu_1+\sigma_1^2}\left(-1+e^{\sigma_1^2}\right)\lambda} + \frac{e^{e^{\mu_1+\frac{\sigma_1^2}{2}}t\lambda}w_{x_1}}{-\beta+e^{\mu_1+\frac{\sigma_1^2}{2}}\lambda}\right) \\
&\quad - e^{-t\beta}\left(\frac{\hat{K}}{\beta} + \frac{e^{t\delta}K_0}{\beta-\delta}\right)w_K\Bigg|_{t=t_0}^{t=T}
\end{aligned} \tag{40}$$

In this decision function, the first two terms represent the discount factors of the system attributes given the threat environment. The third term represents the dynamic evolution of investment in information security as the decision-maker determines the appropriate time horizon, T. In effect, we are mapping the expected loss from the first two terms into the deterministic investment function $K(t)$; see Figure 2. Setting $\mathfrak{D}(t_0, t^*) = 0$ defines the expected investment cycle (the vertical broken line in Figure 2).

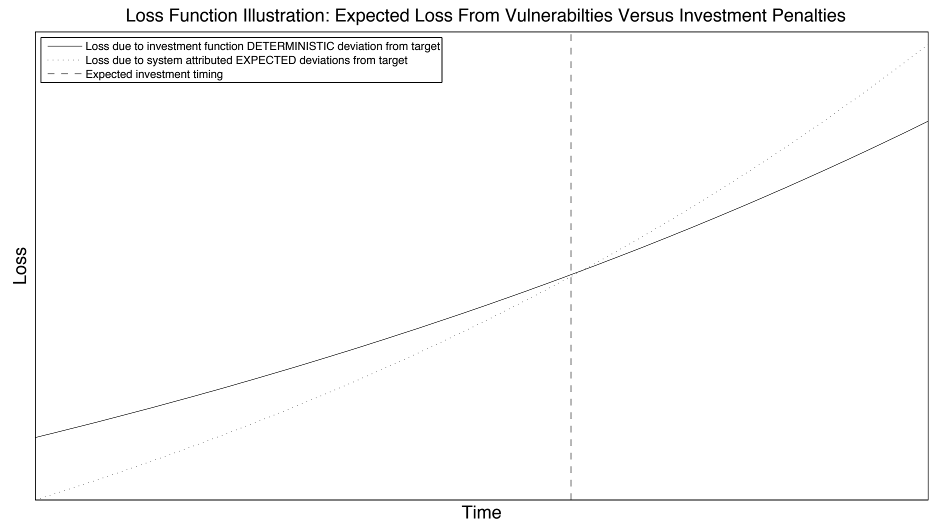

Figure 2: Graphical Illustration of the expected and deterministic losses from deviations from system attributes and investment targets. The rigidity is given by $\hat{K}$, which is the height at which the solid line cuts the loss (vertical) axis.

Thus Proposition 1 describes the equilibrium condition assuming the existence of a solution t^*. The proof of Proposition 1 is outlined in Appendix A.

We now give a characterization of the conditions for the existence of such a solution.

Theorem 1: Existence of Equilibrium

The existence of a solution t^* is subject to the following inequalities obtaining:

$$e^{-t\beta}x_1(t)\left(\frac{e^{e^{2\mu_1+\sigma_1^2}\left(-1+e^{\sigma_1^2}\right)t\lambda}v_{x_1}}{-\beta+e^{2\mu_1+\sigma_1^2}\left(-1+e^{\sigma_1^2}\right)\lambda}+\frac{e^{e^{\mu_1+\frac{\sigma_1^2}{2}}t\lambda}w_{x_1}}{-\beta+e^{\mu_1+\frac{\sigma_1^2}{2}}\lambda}\right)>e^{-t\beta}\left(\frac{\hat{K}}{\beta}+\frac{e^{t\delta}K_0}{\beta-\delta}\right)w_K\quad\forall t>0$$

$$e^{-t\beta}x_2(t)\left(\frac{e^{e^{2\mu_2+\sigma_2^2}\left(-1+e^{\sigma_2^2}\right)t\lambda}v_{x_2}}{-\beta+e^{2\mu_2+\sigma_2^2}\left(-1+e^{\sigma_2^2}\right)\lambda}+\frac{e^{e^{\mu_2+\frac{\sigma_2^2}{2}}t\lambda}w_{x_2}}{-\beta+e^{\mu_2+\frac{\sigma_2^2}{2}}\lambda}\right)>e^{-t\beta}\left(\frac{\hat{K}}{\beta}+\frac{e^{t\delta}K_0}{\beta-\delta}\right)w_K\quad\forall t>0$$

For parsimony, several inequality and boundary constraints are directly inferred from the model structure. The global discount factors, $\delta > 0$, $\beta > 0$, are assumed positive and unbounded. The derivatives are assumed to be in the risk aversion domain, and therefore $w_{x_1}, w_{x_2}, w_K, v_{x_1}, v_{x_2} > 0$, and the initial expected investment point is assumed positive and unbounded; that is, $\hat{K} > 0$.

The proof of Theorem 1 is presented in Appendix B.

4.1 Alternative Cases Under Equilibrium

For the decision function $\mathfrak{D}(\cdot)$, the inequalities presented in Theorem 1 provide three discrete cases:

- if, for all t, the decision function is strictly positive, then the policy-maker addresses vulnerabilties on arrival;
- if, for all t, the decision function is strictly negative, then the policy-maker, in expectations, never anticipates deviation from long-run investment;

- if there is a crossing solution for a positive t, then, in expectations relative to rest, there is an expected investment timing; that is, an investment cycle.

4.2 Discussion

For practical purposes, orientating the security management problem as suggested in Proposition 1 and Theorem 1 has been a useful exercise. Proposition 1 allows us to build a model of timing, by ascribing the system parameters exogenously from data and mapping preferences via a case-study approach to hypothesized events, in the manner of [22]. In related work, to be reported elsewhere, we have used a simulation approach along the lines of [12] and [9], whereby the system architecture of the candidate organization is modelled explicitly, using a location–resource–process systems model. The stochastic properties of the system's attributes are then mapped to our simpler stochastic process and the investment cycle timing.

Subsequently, we can orient the problem using actual data from the organisation on security threats (e.g., using a library such as the NISC CVSS[4] dataset) to populate the inequalities suggested in Theorem 1. We can then orient the free parameters of the system, to see which domains are feasible (i.e., calibrate the stated preferences in the systems model).

In organizations for which confidentiality and availability are of the utmost importance, that is for which the parameters w_a, w_{x_1}, etc. have high values compared to w_K, investment cycles in security will occur relatively more frequently compared to organizations for which investment in information security has high opportunity cost.

For example, in the context of the example of patching policy, a state organization which places very high value on confidentiality, whilst facing soft resource constraints as represented by a low value of w_K, will patch on arrival. A small firm, however, that is severely resource constrained, will avoid undertaking additional investment in information security, that is above $\hat{K}$, to shorten the window of vulnerability. In the strictly negative case of Section 4.1, the model indicates indefinite waiting for investment in the presence of vulnerability threats.

5 Conclusion

In this paper, we have derived (a quadratic approximation to) an analytic solution to the problem of optimal timing of security investments in the presence of existing and future threats. A key aspect of our approach is to introduce the concept of a security investment cycle that is analogous to the classical Keynesian treatment of rigidities.

An interesting potential application of this approach could be in cloud computing, whereby one of the potential macroeconomic benefits is in reducing the amount of capital (fixed) investment. In general, this should reduce vulnerability to negative shocks and, in effect, mitigate business cycles. However, rigidities in security costs could reduce the impact of this effect.

References

[1] R. Abel. Asset Prices under Habit Formation and Catching up with the Joneses. *The American Economic Review*, 80(2):38–42, 1990.

[2] R. Anderson. Why information security is hard: An economic perspective. In *Proc. 17th Annual Computer Security Applications Conference*, pages 358–265. IEEE, 2001.

[3] 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.

[4] National Institute of Science and Technology (NIST), `www.nist.gov`; Common Vulnerability Scoring System (CVSS).

[4] R. Anderson and T. Moore. The economics of information security. *Science*, 314:610–613, 2006. Extended version available at `http://www.cl.cam.ac.uk/~rja14/Papers/toulouse-summary.pdf`.

[5] A. Arora, R. Telang, and H. Xu. Optimal Policy for Software Vulnerability Disclosure. *Management Science*, 54(4):642–656, 2008.

[6] K.J. Arrow. The theory of risk aversion. In *Essays in the Theory of Risk Bearing*, pages 90–109. Markham Publ. Co., 1971. Reprinted from: Aspects of the Theory of Risk Bearing, by Yrjo Jahnssonin Saatio, Helsinki, 1965.

[7] T. August and T. Tunca. Network Software Security and User Incentives. *Management Science*, 52(11):1703–1720, 2006.

[8] Y. Beres, J. Griffin, S. Shiu, M. Heitman, D. Markle, and P. Ventura. Analysing the performance of security solutions to reduce vulnerability exposure window. In *Proceedings of the 2008 Annual Computer Security Applications Conference*, pages 33–42. IEEE Computer Society Conference Publishing Services (CPS), 2008.

[9] Y. Beres, D. Pym, and S. Shiu. Decision support for systems security investment. In *Network Operations and Management Symposium Workshops (NOMS Wksps), 2010 IEEE/IFIP*, pages 118–125, 2010. doi: 10.1109/NOMSW.2010.5486590, ISBN: 978-1-4244-6037-3, INSPEC Accession Number: 11502735.

[10] Nicholas Bloom. The impact of uncertainty shocks. *Econometrica*, 77(3):623–685, 2009.

[11] H. Cavusoglu, H. Cavusoglu, and J. Zhang. Security Patch Management: Share the Burden or Share the Damage. *Management Science*, 54(4):657–670, 2008.

[12] M. Collinson, B. Monahan, and D. Pym. 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.

[13] Larry G. Epstein and Stanley E. Zin. Substitution, Risk Aversion, and the Temporal Behavior of Consumption Growth and Asset Returns I: A Theoretical Framework. *Econometrica*, 57(4):937–969, July 1989.

[14] Peter C. Fishburn. *Utility Theory for Decision Making*. John Wiley and Sons, 1970.

[15] Neal Fultz and Jens Grossklags. Blue versus red: Towards a model of distributed security attacks. In Roger Dingledine and Philippe Golle, editors, *Proceedings of the Thirteenth International Conference Financial Cryptography and Data Security (FC'09)*, pages 167–183. Springer Verlag, February 2009. LNCS 5628, ISBN: 978-3-642-03548-7.

[16] L. Gordon, M. Loeb, and W. Lucyshyn. Information security expenditures and real options: A wait-and-see approach. *Computer Security Journal*, 19(2):1–7, 2003.

[17] 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.

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

[19] C. Ioannidis, D. Pym, and J. Williams. Information security trade-offs and optimal patching policies. Manuscript: Submitted to a journal, 2009.

[20] C. Ioannidis, D. Pym, and J. Williams. Investments and trade-offs in the economics of information security. In Roger Dingledine and Philippe Golle, editors, *Proc. Financial Cryptography and Data Security '09*, volume 5628 of *LNCS*, pages 148–166. Springer, 2009. Preprint available at `http://www.cs.bath.ac.uk/~pym/IoannidisPymWilliams-FC09.pdf`.

[21] D. Kahneman and A. Tversky. Prospect theory: An analysis of decisions under risk. *Econometrica*, 47:313–327, 1979.

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

[23] Otto 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.

[24] M. Casassa Mont, Y. Beres, 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*, pages 134–141, 2010. doi: 10.1109/NOMSW.2010.5486588, ISBN: 978-1-4244-6037-3, INSPEC Accession Number: 11502733.

[25] J.W. Pratt. Risk aversion in the small and in the large. *Econometrica*, 32:122–136, January–April 1964.

[26] D. Rogers and L.C.G. Williams. *Diffusions, Markov Processes, and Martingales.* Cambridge Mathematics Library, 2000.

[27] S.M. Ross. *Stochastic Processes.* Wiley, 1995.

[28] K. Taksumi and M. Goto. Optimal timing of information security investment: A real options approach. In T. Moore, D. Pym, and C. Ioannidis, editors, *Economics of Information Security and Privacy.* Springer, 2010. Proceedings WEIS 2009, London.

A Proof of Proposition 1: Decision Function Equilibrium

The vector random variable z has n univariate log-normal marginal distributions and correlation structure driven by $\mathbf{\Sigma} = \mathbb{E}\left(\log z - \mu\right)\left(\log z - \mu\right)'$. Where $'$ denotes the conjugate transpose, $\mu = \{\mu_1, \ldots, \mu_n\}$ is a vector of time-homogenous central expectations, and $\mathbf{\Sigma} = [\sigma_{ij}]$ is a time-homogenous covariance variance matrix. Setting $n = 2$, the observed moments and co-moments of z are defined as

$$\mathbb{E}\left(z_1\right) = e^{\mu_1 + \frac{\sigma_1^2}{2}} \tag{41}$$

$$\mathbb{E}\left(z_2\right) = e^{\mu_2 + \frac{\sigma_2^2}{2}} \tag{42}$$

$$\mathbb{E}\left(z_1 - \mathbb{E}\left(z_1\right)\right)^2 = \operatorname{var}\left(z_1\right) = e^{2\mu_1 + \sigma_1^2}\left(-1 + e^{\sigma_1^2}\right) \tag{43}$$

$$\mathbb{E}\left(z_2 - \mathbb{E}\left(z_2\right)\right)^2 = \operatorname{var}\left(z_1\right) = e^{2\mu_2 + \sigma_2^2}\left(-1 + e^{\sigma_2^2}\right) \tag{44}$$

$$\mathbb{E}\left(z_1 - \mathbb{E}\left(z_1\right)\right)\left(z_2 - \mathbb{E}\left(z_2\right)\right) = \operatorname{cov}\left(z_1, z_2\right) = e^{\mu_1 + \mu_2 + \frac{1}{2}\left(\sigma_1^2 + \sigma_2^2\right)}\left(-1 + e^{\sigma_{12}}\right) \tag{45}$$

where μ_1 is the expected value of the generating normal distribution for the system attribute x_1, μ_2 is the expected value of the generating normal distribution for the system attribute x_2, σ_1 is the standard deviation of the underlying normal distribution for x_1, σ_2 is the standard deviation of the underlying normal distribution for x_2, and σ_{12} is the covariance of the underlying normal. distribution.

Consider an arrival rate λ. The expected number of events over the interval $[t_0, T]$ is then $\lambda(T - t_0)$. Let $t_0 \leq t \leq T$ and set $t_0 = 0$. The combined jump process $y(t) \in \mathbb{R}^2$, with poisson arrivals with count $\lambda(t)$, yields the expected moments of the threat environment discount factors for the interval $[0,t]$ as follows:

$$
\begin{aligned}
\mathbb{E}\left(y_1\right) &= e^{\mu_1+\frac{\sigma_1^2}{2}}\lambda\left(t\right) && (46)\\
\mathbb{E}\left(y_2\right) &= e^{\mu_2+\frac{\sigma_2^2}{2}}\lambda\left(t\right) && (47)\\
\mathbb{E}\left(y_1-\mathbb{E}\left(y_1\right)\right)^2 &= \operatorname{var}\left(y_1\right) = e^{2\mu_1+\sigma_1^2}\left(-1+e^{\sigma_1^2}\right)\lambda\left(t\right) && (48)\\
\mathbb{E}\left(y_2-\mathbb{E}\left(y_2\right)\right)^2 &= \operatorname{var}\left(y_1\right) = e^{2\mu_2+\sigma_2^2}\left(-1+e^{\sigma_2^2}\right)\lambda\left(t\right) && (49)\\
\mathbb{E}\left(y_1-\mathbb{E}\left(y_1\right)\right)\left(y_2-\mathbb{E}\left(y_2\right)\right) &= \operatorname{cov}\left(y_1,y_2\right) = e^{\mu_1+\mu_2+\frac{1}{2}\left(\sigma_1^2+\sigma_2^2\right)}\left(-1+e^{\sigma_{12}}\right)\lambda\left(t\right) && (50)
\end{aligned}
$$

Combining the moments of the system process from Equation 39 with the instantaneous expectations from Equations (14-18), yields

$$
\begin{aligned}
\mu_{x_1}\left(t\right) &= e^{e^{\mu_1+\frac{\sigma_1^2}{2}}(t)\lambda}x_1(t) && (51)\\
\mu_{x_2}\left(t\right) &= e^{e^{\mu_2+\frac{\sigma_2^2}{2}}(t)\lambda}x_2(t) && (52)\\
\sigma_{x_1}\left(t\right) &= e^{e^{2\mu_1+\sigma_1^2}\left(-1+e^{\sigma_1^2}\right)(t)\lambda}x_1(t) && (53)\\
\sigma_{x_2}\left(t\right) &= e^{e^{2\mu_2+\sigma_2^2}\left(-1+e^{\sigma_2^2}\right)(t)\lambda}x_2(t) && (54)\\
\sigma_{x_1,x_2}\left(t\right) &= e^{e^{\mu_1+\mu_2+\frac{1}{2}\left(\sigma_1^2+\sigma_2^2\right)}(-1+e^{\sigma_{12}})(t)\lambda}x_2(t)x_1(t) && (55)
\end{aligned}
$$

Substituting these into Equation (30) gives the explicit instantaneous loss function, for compactness we assume separable additivity, therefore all the covariance terms drop out as follows:

$$
\begin{aligned}
\ell\left(t|w_{x_1},w_{x_2},v_{x_1},v_{x_2},v_{x_1,x_2}\right) = \; & e^{e^{2\mu_2+\sigma_2^2}\left(-1+e^{\sigma_2^2}\right)(t)\lambda}x_2(t)v_{x_2} + \\
& e^{e^{2\mu_1+\sigma_1^2}\left(-1+e^{\sigma_1^2}\right)(t)\lambda}x_1(t)v_{x_1} + \\
& e^{e^{\mu_1+\mu_2+\frac{1}{2}\left(\sigma_1^2+\sigma_2^2\right)}(-1+e^{\sigma_{12}})(t)\lambda}x_2(t)x_1(t)v_{x_1,x_2} + \\
& e^{e^{\mu_2+\frac{\sigma_2^2}{2}}(t)\lambda}x_2(t)w_{x_2} + \\
& e^{e^{\mu_1+\frac{\sigma_1^2}{2}}(t)\lambda}x_1(t)w_{x_1} \qquad (56)
\end{aligned}
$$

Integrating and discounting over the policy-maker's time horizon yields the expected loss function at time t_0,

$$
\begin{aligned}
\mathfrak{L}\left(t_0,T|w_{x_1},w_{x_2},v_{x_1},v_{x_2},v_{x_1,x_2}\right) = \; & -e^{-t\beta}x_2(t)\frac{e^{-e^{2\mu_2+\sigma_2^2}\left(-1+e^{\sigma_2^2}\right)(t)\lambda}v_{x_2}}{\beta+e^{2\mu_2+\sigma_2^2}\left(-1+e^{\sigma_2^2}\right)\lambda} \\
& -e^{-t\beta}x_2(t)\frac{e^{-e^{\mu_1+\mu_2+\frac{\sigma_1^2}{2}+\frac{\sigma_2^2}{2}}(-1+e^{\sigma_{12}})(t)\lambda}x_1(t)v_{x_1,x_2}}{\beta+e^{\mu_1+\mu_2+\frac{\sigma_1^2}{2}+\frac{\sigma_2^2}{2}}\left(-1+e^{\sigma_{12}}\right)\lambda} \\
& -e^{-t\beta}x_2(t)\frac{e^{-e^{\mu_2+\frac{\sigma_2^2}{2}}(t)\lambda}w_{x_2}}{\beta+e^{\mu_2+\frac{\sigma_2^2}{2}}\lambda} - e^{-t\beta}x_1(t)\frac{e^{-e^{2\mu_1+\sigma_1^2}\left(-1+e^{\sigma_1^2}\right)(t)\lambda}v_{x_1}}{\beta+e^{2\mu_1+\sigma_1^2}\left(-1+e^{\sigma_1^2}\right)\lambda} \\
& \left. -e^{-t\beta}x_1(t)\frac{e^{-e^{\mu_1+\frac{\sigma_1^2}{2}}(t)\lambda}w_{x_1}}{\beta+e^{\mu_1+\frac{\sigma_1^2}{2}}\lambda}\right|_{t=t_0}^{t=T} \qquad (57)
\end{aligned}
$$

B Proof of Theorem 1

Let $T < \infty$. For $t^* \in (t_0, T)$, the cost function at t_0 is defined as

$$\begin{aligned}
\mathfrak{U}(t_0, T|w_k) &= \int_{t_0}^{T} e^{-\beta t}\left(\hat{K} + e^{t\delta}K_0\right) w_k dt \\
&= e^{-t\beta}\left(-\frac{\hat{K}}{\beta} + \frac{e^{t\delta}K_0}{-\beta+\delta}\right) w_k \Bigg|_{t=t_0}^{t=T} \qquad (58)
\end{aligned}$$

The subsequent trade-off decision-making function relative to the cost function from 35 is now

$$\begin{aligned}
\mathfrak{D}(t_0, T|w_{x_1}, w_{x_2}, w_K, v_{x_1}, v_{x_2}, v_{x_1,x_2}) &= -e^{-t\beta} x_2(t) \frac{e^{-e^{2\mu_2+\sigma_2^2}\left(-1+e^{\sigma_2^2}\right)\lambda(t)} v_{x_2}}{\beta + e^{2\mu_2+\sigma_2^2}\left(-1+e^{\sigma_2^2}\right)\lambda} \\
&\quad -e^{-t\beta} x_2(t) \frac{e^{-e^{\mu_1+\mu_2+\frac{\sigma_1^2}{2}+\frac{\sigma_2^2}{2}}(-1+e^{\sigma_{12}})\lambda(t)} x_1(t) v_{x_1,x_2}}{\beta + e^{\mu_1+\mu_2+\frac{\sigma_1^2}{2}+\frac{\sigma_2^2}{2}}(-1+e^{\sigma_{12}})\lambda} \\
&\quad -e^{-t\beta} x_2(t) \frac{e^{-e^{\mu_2+\frac{\sigma_2^2}{2}}\lambda(t)} w_{x_2}}{\beta + e^{\mu_2+\frac{\sigma_2^2}{2}}\lambda} \\
&\quad +x_1(t)\left(-\frac{e^{-e^{2\mu_1+\sigma_1^2}\left(-1+e^{\sigma_1^2}\right)\lambda(t)} v_{x_1}}{\beta + e^{2\mu_1+\sigma_1^2}\left(-1+e^{\sigma_1^2}\right)\lambda} - \frac{e^{-e^{\mu_1+\frac{\sigma_1^2}{2}}\lambda(t)} w_{x_1}}{\beta + e^{\mu_1+\frac{\sigma_1^2}{2}}\lambda}\right) \\
&\quad + \left(\frac{\hat{K}}{\beta} + \frac{e^{t\delta}K_0}{\beta-\delta}\right) w_k \Bigg|_{t=t_0}^{t=T} \qquad (59)
\end{aligned}$$

It is now trivial to see that the inequalities in Theorem 1 are obtained by subtracting last term of Equation 59 from each side and setting either $x_1(t_0)$ or $x_2(t_0)$ to zero.