

Is Public Co-Ordination of Investment in Information Security Desirable?

Christos Ioannidis¹, David Pym², Julian Williams³

¹Department of Economics, University of Bath, Bath, UK

²Department of Computer Science, University College London, London, UK

³Durham University Business School, Durham, UK

Email: c.ioannidis@bath.ac.uk, d.pym@ucl.ac.uk, julian.williams@durham.ac.uk

Received 21 March 2016; accepted 27 March 2016; published 30 March 2016

Copyright © 2016 by authors and Scientific Research Publishing Inc.

This work is licensed under the Creative Commons Attribution International License (CC BY).

<http://creativecommons.org/licenses/by/4.0/>



Open Access

Abstract

This paper provides for the presentation, in an integrated manner, of a sequence of results addressing the consequences of the presence of an information steward in an ecosystem under attack and establishes the appropriate defensive investment responses, thus allowing for a cohesive understanding of the nature of the information steward in a variety of attack contexts. We determine the level of investment in information security and attacking intensity when agents react in a non-coordinated manner and compare them to the case of the system's coordinated response undertaken under the guidance of a steward. We show that only in the most well-designed institutional set-up the presence of the well-informed steward provides for an increase of the system's resilience to attacks. In the case in which both the information available to the steward and its policy instruments are curtailed, coordinated policy responses yield no additional benefits to individual agents and in some case they actually compared unfavourably to atomistic responses. The system's sustainability does improve in the presence of a steward, which deters attackers and reduces the numbers and intensity of attacks. In most cases, the resulting investment expenditure undertaken by the agents in the ecosystem exceeds its Pareto efficient magnitude.

Keywords

Information Security, Information Stewardship, Investment, Public Co-Ordination

1. Information Stewardship

Information produces value for an organization or individual when it improves the solutions to decision-making problems whose outcomes have consequences for their welfare. The information system refers to the entire

collection of data sources and related service capabilities both internal and external to the organization that decision makers are required to use. The system is user-centred because it serves the objectives of the organization by providing the information needed to achieve its mission.

The security of such systems is of paramount importance and economic agents are willing to allocate scarce resources in protecting the system when it is threatened. Information systems based almost exclusively on digital technologies are subject to and degraded by cyberattacks, which are initiated and executed remotely. As the subjects of such attacks have no clear means of identifying the initiators and stopping their activities, their main concern is to preserve the system's functionality in all its dimensions by allocating resources, thus incurring cost, to maintain it at the level of operational capacity required by the organization. More specifically, information security, is conventionally defined as protecting the system's confidentiality, integrity and availability (CIA).

The aim of this paper is to examine whether the decisions about expenditure/investment in information security should be socially coordinated via a *steward* or such decisions are better left to individual organizations. The subsequent discussion and models follow closely [1]-[4], where the proofs elided here may be found. This paper's contribution lies in the integration of a dispersed body of work addressing the issues regarding the co-ordination of investment in information security. Whilst the problems of sustainability and resilience have been addressed separately, combining them in a single theoretical framework provides for a clear appreciation of the policy issues emerging from public co-ordination. Related organizational issues have been examined by [5]. There is no generally accepted scientific definition of the concept of stewardship. In more general terms, stewardship is an ethic guiding the allocation and management of some of the participants' resources in an ecosystem (household, common interest community, commercial firm, etc.) in order to sustain and protect the ecosystem, rather than the welfare maximization of individual agents, in the presence of anticipated and unanticipated shocks. The steward will be part of the ecosystem itself and can emerge either from internal or external forces.

The concept of stewardship in environmental economics is an established tool for environmental and natural resource management; see [6] and mitigation risk for climate change [7]. In the context of information security, we define the role of the steward as the institution that maintains the sustainability and resilience of the ecosystem's operational capacity—that is, its levels of confidentiality, integrity, and availability—which may be threatened by attacks. By resilience (cf. [8]), we mean the system's internal capacity to restore itself to an acceptable operating state following a disturbance to its status; by sustainability, we mean the tendency of the system to maintain itself within acceptable bounds of operating state despite possibly hidden dynamics that may tend to guide the system outwith these bounds.

Stewards might emerge as a consequence of the behaviour of agents interacting in a system of exchange which is in turn conditioned by their preferences, the established legal framework, and existing social conventions. Such conventions, known as norms, are either descriptive—that is, what actions the agents in the system take—or prescriptive, influencing what behaviour ought to be. The legal framework expresses system's values and determines the consequences (punishment) for actions deviating from such values. [9] argued that agents derive benefits from the supply of the public good—in this context, sustainability and resilience for the ecosystem—and, more importantly, that they have an intrinsic motivation to undertake costly effort to the production of the public good.

Part of the role of the steward is to alter what constitutes normal behaviour. Agents in a decentralized ecosystem may have beliefs, based on incomplete information, regarding the contributions of the others and thus are in danger in holding erroneous perceptions of the true societal norm. The steward, by dispelling such miss-perceptions, can attain substantial benefits for the system as participants modify their behaviour under its guidance. For example, when agents are excessively optimistic regarding the conduct of others, the result is a fall in compliance. In this context, the steward of the ecosystem, which is subject to shocks and in secular decline, to protect the system will exercise the option available to him, prescriptive intervention. Such interventions range from widely publicized public campaigns to enforceable standards, and which boost social pressure on the individual agents to comply and make punishable the failure by agents to meet these standards. The underlying assumption here is that the steward is globally-informed compared to individuals about the currently prevailing community standards. In a more general setting, the steward knows the underlying distribution of preferences/risk aversion/discount rates in society, information which is difficult and costly for a single agent to collect and process. The importance of the legal framework in delivering binding agreement to the production of the public good has been studied by [10] [11], and others. Such studies show that compliance is raised when its level has been cho-

sen through a voting decision by the participants of the community/ecosystem. To maximize the effectiveness of its actions, the law-maker/principal/steward, in setting the regulatory framework and other obligations/incentives, must take into account the impact of its actions on the formulation of the norms which will be now expected to prevail. To illustrate the point in the context of investment in cybersecurity we may consider a steward signaling to its community the need for very high levels of IT defence expenditure. Such request is actually conveying the signal that the current situation is very dangerous and in this case this might deter well meaning agents as they perceive themselves as spending too much compared to their community. It is important that the steward allows a framework where the behaviour of the individual is observed by others to ensure compliance to the chosen standard.

The exercise of stewardship is costly to the ecosystem participants and to agents assuming this role as it requires investment of resources in infrastructure, the creation of a legal framework and monitoring required to deliver its mission. The motives for agents to engage in such activities are diverse, and they range from individual welfare maximization to a form of altruistic behaviour. Simple profit maximization can be seen as the motive of a retailer, who voluntarily assumes the role of the steward, in the case where multiple sellers are using its platform to connect with consumers. In this case the retailer acting as the steward safeguards that contracts negotiated between the sellers and the consumers on its platform are honoured and the ecosystem maintains its credibility as a space of secure commercial transactions.

Profit maximization is by no means the only motive for engaging in stewardship. Recent studies from psychology and economics provide strong empirical evidence for the existence of pro-social behaviour. For a recent survey, see [9] [12] [13]. That is, agents engage frequently in costly activities whose benefits accrue to others. In a highly decentralized ecosystem, some agents will possess the intrinsic motivation to behave pro-socially. Such agents will typically have two motivations. First, they will care for the overall provision of the public good to which their individual actions contribute, but also they will care for the consumption of this good by others.

We begin this paper by developing separate models, to elucidate in some detail the concepts of resilience and sustainability in terms of investment in information security with and without the information steward. Both models presented in Sections 2 and 3 (following [1]-[4]) are based on strategic interactions between the agents of the eco-system who wish to protect their information resources and attackers who derive benefits from degrading the system's operational capacity. Both attackers and defenders are incurring costs whilst engaging in their respective activities and the extent of their involvement in defending and attacking is determined by the existing technologies, time preferences and the value of the information assets.

The two graphs below depict the concepts in terms of the time evolution of the system's operational capacity in the presence, first, of secular deterioration, and, second, unanticipated performance reducing shocks. The dynamics of resilience are depicted in **Figure 1**. In this graph, we depict the system's predictable time path within the acceptable tolerances in performance denoted by its nominal operational capacity. Along this path, at time t_1 the system experiences an unanticipated shock of moderate magnitude that degrades its capacity, placing it outside the acceptable range and guides the system to lower capacity levels. In the absence of the steward, such a shock may prove permanently detrimental to the state of the system, with the system's path depicted by the broken line. The actions of the steward render the system resilient as they are able to reverse the divergent path and restore the system to its "trend" capacity (solid line), up to the planning horizon, T . Alternatively, in the presence of a substantial shock, as the one depicted at time t_2 , the best the steward can achieve is to halt the system's rapid deterioration and stabilize the system's operational capacity to a steady, albeit lower, level.

The dynamics of sustainability are depicted in **Figure 2**. In this graph, we characterize the system's equilibrium course over time. We envisage that the system degrades steadily and predictably along this path. Its internal dynamic structure without the steward will result in the system's inability to perform within the acceptable bounds by t_1 . The steward's contribution to the system's sustainability is to delay the rate of degradation beyond the planning horizon T . Again, the steward adopts the relevant policies and installs the required institutional framework at t_0 . Therefore, the steward permanently changes the long-term dynamic structure of the system at the beginning, permitting the system to enjoy a considerable extension to its useful life compared to the state where the steward is absent.

By adopting such actions/policies, the steward extends the life of the system, minimizes the impact of the shocks, enhancing the system's predictability and robustness of performance. The benefits of the existence of such institutional arrangement accrue to all the participants of the ecosystem, reducing the incentive for existing members to abandon and encouraging new agents to join. **Figure 3** nests the previous notions of stewardship

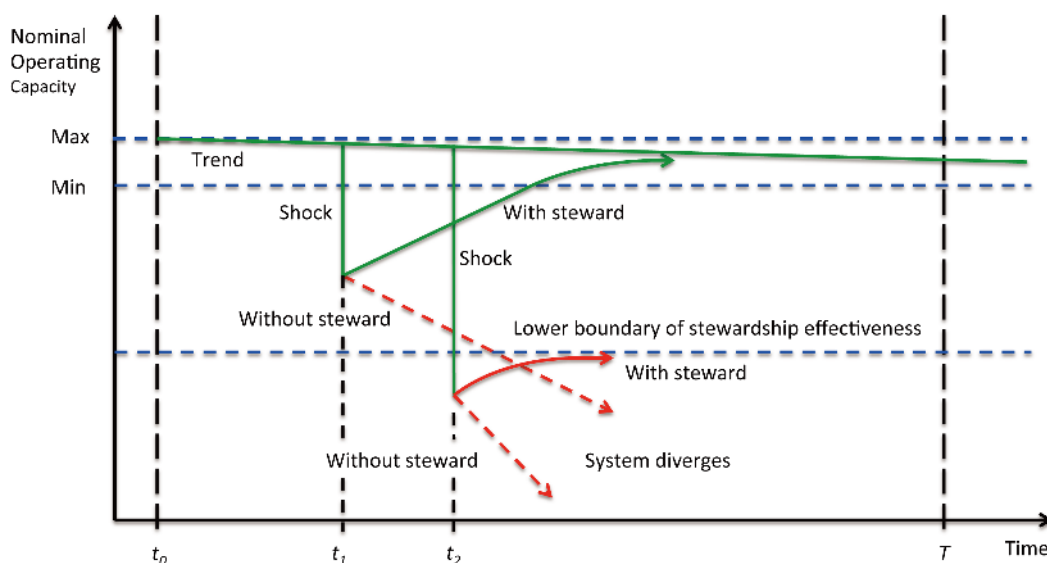


Figure 1. Resilience.

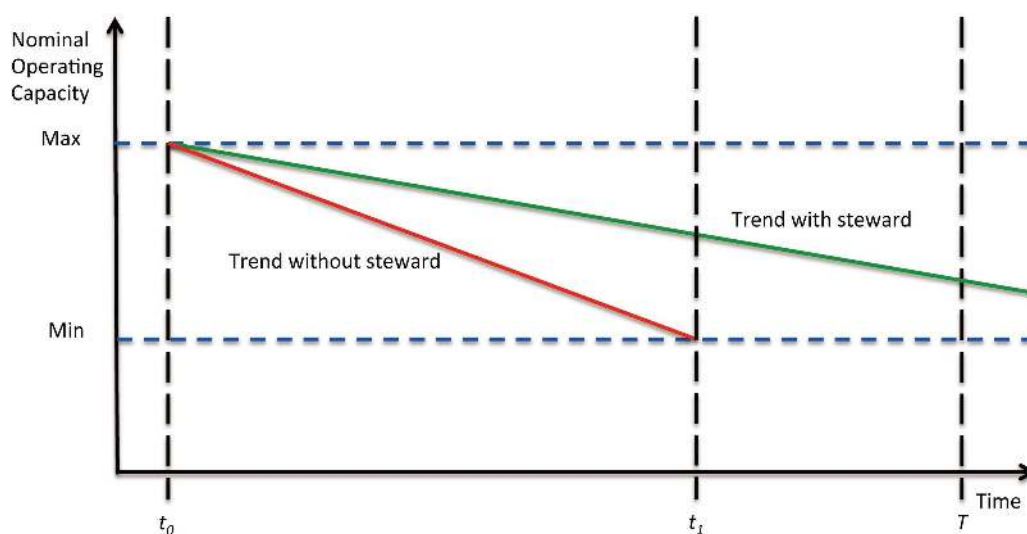


Figure 2. Sustainability.

into a multi-period representation of system sustainability.

There are prominent examples of the exercise of information stewardship in EU, UK, and US legislation. For example, consider the contribution of the existence of the Freedom of Information Acts in the UK and US on policy decision-making. In their absence, information provided by the public will be limited by the public's perception of misuse, so restricting the information available to policy-makers. For another example, the US's response to EU privacy legislation in constituting "Safe Harbor" (<http://export.gov/safeharbor/>, access 4 March 2013) encouraged and maintained trade between the two economics. Both cases are examples of sustainability as the steward intervenes to maintain the market.

The failure of stewardship to maintain resilience is demonstrated by the failure in June 2012 of the Royal Bank of Scotland's payment processing systems, which support an ecosystem subsidiary banks. A software upgrade corrupted the system and, in the absence of sufficient system management resources, the ecosystem's payment processing systems ceased to function for a considerable period of time and, for several banks, acceptable service levels were restored only after considerable delay.

The remainder of this paper is organized as follows: Section 2 provides a summary of mathematical results on

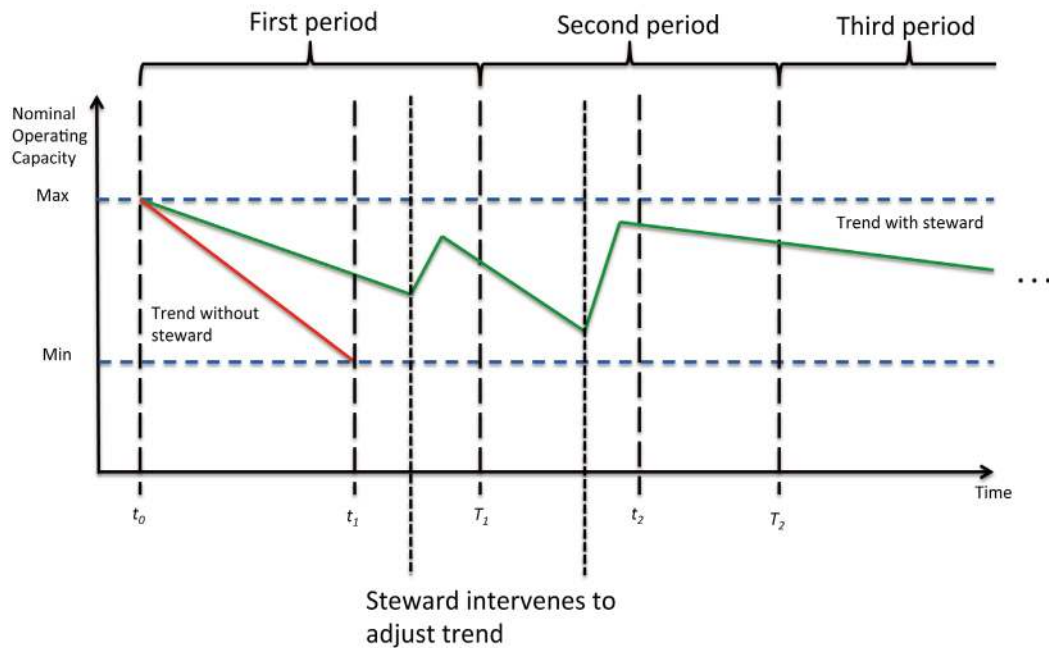


Figure 3. Nesting one period sustainability into a multi-period framework.

the notion of sustainability, and the proofs for this section may be found in [2]; Section 3 provides a summary of a mathematical model of resilience in a simple security setting, and the proofs for this section may be found in [4]; finally, Section 4 provides a brief summary of the implications for stewardship in this setting.

2. A Model of Sustainability

We begin with a condensed outline of the model presented in [1]-[4].

We assume that the steward is able to pass on the cost of his decisions to the agents in the ecosystem whilst improving the agents' security. In this set-up, individual agents are assumed face a known probability of successful attack and they possess known mitigation technologies. In both cases, to demonstrate the role of a steward in enhancing the ecosystem's sustainability and resilience, we set up a model of strategic interactions between targets (typically firms, but possibly individuals) and attackers (typically individuals, but possibly organized groups). Attackers are assumed to be profit-maximizing and risk-neutral.¹ We introduce the concept of a steward who co-ordinates the defensive expenditure of the targets. We show that as the steward, as in **Figure 2**, seeks to sustain the life of the ecosystem to at least the planning horizon is setting mandatory levels of defensive expenditure for each target. These are set using time preferences for the valuation of expected future losses from successful attacks over the period of his planning horizon.

We begin by evaluating the role of the steward in contributing to the system's sustainability. We assume that the steward is a Stackelberg policy-maker, who imposes a policy as a first move irrespective of the reactions of the participants in the ecosystem. In addition, we assume that increasing defensive expenditure by an ecosystem participant reduces the likelihood of successful attacks on that participant and that for any given participant, the greater the number of attackers, the higher the likelihood of a successful attack.

Building on the seminal contribution to benefit-cost analysis for information security investments presented by Gordon and Loeb (GL) [14], we present a framework for managing security investments in information ecosystems [5] in which we identify the role of the steward in regulating the allocation of resources by the ecosystem's participants.

Targets of security attacks are GL expected-loss-minimizers and Attackers act as rational agents. They are assumed to have utility functions, with well-defined preferences. All attack and investment decisions are taken

¹It is straightforward to vary the model to account for payoffs that are utility-maximizing, but not necessarily monetary; for example, as in terrorism.

discounting under a risk-neutral measure. That is, the combination of time preferences (captured by discount factors, β), probability measures, $\tilde{\Psi}$, and measured losses, L , admits any standard representation (e.g., constant relative or constant absolute) of risk-aversion. The possible difference between the discount factors, β , adopted by targets and the steward will be of importance in evaluating the stewards's contribution to the system's sustainability.

Consider an ecosystem with N_T targets and a fixed number N_A of attackers. We set the ratio $\eta = \frac{N_A}{N_T}$; that is, the number of attackers per target.

Decisions on security investment are taken at time zero and are assumed to be made with full commitment. The usefulness of time in this context is not to add temporal dynamics to the security investment problem, but to illustrate the impact of different discount rates between participants in the game.

Let $\Psi(\eta, \alpha_i, t)$ be the instantaneous probability that a single attack will be successful in the absence of any defensive expenditure. We will assume that attacks have independent probabilities and a functional form that satisfies the condition that eventually an attacker will be successful is

$$\Psi(\eta, \alpha_i, t) = 1 - e^{-\alpha_i \eta t},$$

where α_i is a technology parameter, *the security decay factor*, that relates the probability of successful attack to the number of attackers per target and t is continuous time in the interval $t_0 < t \leq T$.

The expected loss at t , in the absence of security investment, is given by $L_i \Psi(\eta, \alpha_i, t) \equiv L_i (1 - e^{-\alpha_i \eta t})$, where L_i is the current or nominal monetary loss from an attack and is fixed over $t_0 < t \leq T$. We assume the current value of assets is evenly (by time discounted weighting) amortized over T .

The present value of losses, in the absence of security investment, given a discount rate β_i , for the i^{th} target is

$$\int_{t_0}^T e^{-\beta_i t} L_i (1 - e^{-\alpha_i \eta t}) dt = \left[-\frac{L_i e^{-t(\alpha_i \eta + \beta_i)} (\beta_i (e^{\alpha_i \eta t} - 1) + \alpha_i \eta e^{\alpha_i \eta t})}{\beta_i (\alpha_i \eta + \beta_i)} \right]_{t_0}^T.$$

Let $\tilde{\Psi} = \Psi(\eta, \alpha_i, t) e^{-\psi_i x}$. $\tilde{\Psi}$ is the instantaneous probability of realizing a loss L_i and the expected value of losses over the time period $[t_0, T]$ is given by

$$\int_{t_0}^T e^{-\beta_i t} \tilde{\Psi} L_i dt.$$

That is, integrating losses multiplied by their probabilities and discounted at rate β .

2.1. Target Security Investment

As targets are aware of the threats to their information assets each target has a control instrument, denoted x_i , the level of defensive expenditure and, for simplicity of exposition, we also assume that this is set at t_0 with commitment.

We further assume that defensive expenditure reduces the probability of a successful attack by a continuous rate, ψ_i , which is another technology parameter, *the security effectiveness factor*. The interpretation of ψ_i is that it is the amount of investment needed to reduce the probability of attack by $1/e$, following the Gordon and Loeb rule [14].

Therefore, in the presence of defensive expenditure the instantaneous expected loss from attacks, in the presence of defensive expenditure at time t , is now $L_i (1 - e^{-\alpha_i \eta t}) e^{-\psi_i x_i}$. Setting $\psi_i = 0$, makes instantaneous expected losses constant and independent of defensive expenditure x_i . That is, targets are incapable, for all t , of mitigating the risk of loss.

The term ψ relates the effectiveness of defensive expenditure in mitigating the probability of a successful attack in all periods. Taking into account the investment in information security the expected present value of losses for each target is therefore

$$\begin{aligned}
PV &= \int_{t_0}^T e^{-\beta_i t} L_i (1 - e^{-\alpha_i \eta t}) e^{-\psi_i x_i} dt \\
&= \left[\frac{L_i (\beta_i (e^{\alpha_i \eta T} - 1) + \alpha_i \eta e^{\alpha_i \eta T}) e^{t(-\alpha_i \eta - \beta) - x \psi_i}}{\beta_i (\alpha_i \eta + \beta_i)} \right]_{t_0}^T.
\end{aligned}$$

As targets are risk neutral (relative to their discount rate), the net present value (adding the t_0 expenditure of $-x_0$) of losses is equivalent to their utility at t_0 . Therefore,

$$U(x_i) = -\tilde{L}_i(x_i) = \left[-\frac{L_i (\beta_i (e^{\alpha_i \eta T} - 1) + \alpha_i \eta e^{\alpha_i \eta T}) e^{t(-\alpha_i \eta - \beta) - x \psi_i}}{\beta_i (\alpha_i \eta + \beta_i)} \right]_{t_0}^T - x_i.$$

In the presence of an exogenous η , the i^{th} target minimizes losses with respect to x_i :

$$x_i^* = \operatorname{argmin}_{x_i} \frac{L_i (\beta_i (e^{\alpha_i \eta T} - 1) + \alpha_i \eta e^{\alpha_i \eta T}) e^{T(-\alpha_i \eta - \beta) - x \psi_i}}{\beta_i (\alpha_i \eta + \beta_i)} + x_i.$$

Differentiating the net present value of losses with respect to x_i and setting the derivative equal to zero yields

$$-\frac{L_i \psi_i e^{-T(\alpha_i \eta + \beta_i) - x_i \psi_i}}{\alpha_i \eta + \beta_i} + \frac{L_i \psi_i e^{-T\beta_i - x_i \psi_i}}{\beta_i} = 1.$$

Therefore, for a given η , x^* has the following analytic solution:

$$x_i^*(\eta, \alpha_i, \beta_i, \psi_i, L_i, T) = \psi_i^{-1} \log \left(\beta_i^{-1} L_i \psi_i e^{-\beta_i T} - (\alpha_i \eta + \beta_i)^{-1} L_i \psi_i e^{-T(\beta_i + \alpha_i \eta)} \right).$$

Thus improvements in protective technology (increasing ψ_i s) lead to diminishing optimal marginal security returns on expenditure and, as the efficiency of attackers increases (increasing α_i s), optimal defensive expenditure increases slightly more than proportionately.

For any given monetary loss, L_i , time horizon, T , and technology parameters, ψ_i and α_i , increases in the discount rate, β_i , lead to lower defensive expenditures as the valuation of future losses declines faster.

Attackers are non-cooperative and risk neutral and make rational choices to participate in attacks. The reward for individual attacker successfully attacking agent I is denoted R_i . We further assume that attacking effort requires a costly one-off investment at t_0 , denoted by C_A , and that future gains from attacks are discounted at a rate γ . We suppose that attackers randomly choose targets according to a uniform distribution and as such can only identify average defensive expenditure \tilde{x} where $\tilde{x} = \sum_{i=1}^{N_T} x_i$. The expected reward a time t from successful attacks is given by

$$V(t) = N_A^{-1} \sum_{i=1}^{N_T} R_i (1 - e^{-\tilde{\alpha} t}) e^{-\tilde{\psi} \tilde{x}_i},$$

and its present value by

$$PV_A(t_0, T) = \left[\frac{\tilde{R} (\gamma + \tilde{\alpha} e^{t(\tilde{\alpha} + \gamma)} - (\tilde{\alpha} + \gamma) e^{\tilde{\alpha} T}) e^{-t(\tilde{\alpha} + \gamma) - \tilde{x} \tilde{\psi}}}{\gamma \eta (\tilde{\alpha} + \gamma)} \right]_{t_0}^T.$$

The marginal attacker enters the market until the present value of expected rewards, $PV_A(T)$, equals the value of costs C_A . In this case of a single decision to attack, with first winner takes all attackers and random target selection the attacker decision reduces to the expectation of being the successful attacker from η attackers. The boundary condition for the marginal attacker choosing to enter the market for attacks requires the equality of the present value of an attack PV_A to its cost of launching it:

$$\frac{\tilde{R} \left(\gamma + \tilde{\alpha} e^{T(\tilde{\alpha} + \gamma)} - (\tilde{\alpha} + \gamma) e^{\tilde{\alpha} T} \right) e^{-T(\tilde{\alpha} + \gamma) - \tilde{x}\tilde{\psi}}}{\gamma \eta (\tilde{\alpha} + \gamma)} = C_A.$$

Dividing both sides of this equation by R , setting $\tilde{c} = C_A / \tilde{R}$ to be the expected cost per reward, and solving for η , the equilibrium level of attacks per target is

$$\eta^* = \frac{\left(\gamma + \tilde{\alpha} e^{T(\tilde{\alpha} + \gamma)} - \gamma e^{\tilde{\alpha} T} - \tilde{\alpha} e^{\tilde{\alpha} T} \right) e^{-T(\tilde{\alpha} + \gamma) - \tilde{x}\tilde{\psi}}}{\tilde{c} \gamma (\tilde{\alpha} + \gamma)}.$$

The following proposition, proved in [1] [2], establishes the level of optimal expenditure in information security by each target and the attack intensity (attacks per target):

Proposition 1. For N_T ex-ante identical targets choosing defensive expenditure level x and for N_A first-winner-takes-all attackers, where $\eta = N_A / N_T$ is endogenous, the Nash equilibrium levels of expenditure x^N and the number of attackers per target η^N are given by the solutions to the following pair of equations:

$$\eta^N = c\gamma^{-1} (\alpha + \gamma)^{-1} \left(\gamma + \alpha e^{T(\alpha + \gamma)} - \gamma e^{\alpha T} - \alpha e^{\alpha T} \right) e^{-T(\alpha + \gamma) - x^N \psi}$$

$$x^N = \psi^{-1} \log \left(\beta^{-1} L \psi e^{-\beta T} - (\alpha \eta^N + \beta)^{-1} L \psi e^{-T(\beta + \alpha \eta^N)} \right).$$

When the target responses are coordinated by a fully informed information steward whose sole objective is to improve the ecosystem's sustainability and can impose his choice of defensive expenditure, x_i^P , on each individual target, i , in calculating the optimal investment in cyber-defence takes into account its impact on the actions of attackers.

The steward minimizes his objective function, where the x_i^P denote the investments required by the steward for targets i :

$$\left[x_i^P \right]_{i=1}^{N_T} = \arg \min_{\left[x_i^P \right]_{i=1}^{N_T}} \sum_{i=1}^{N_T} \int_{t_0}^T e^{-\delta_i t} L_i \left(1 - e^{-\alpha_i \eta(\mathbf{x}^P) t} \right) e^{-\psi_i x_i} dt + \sum_{i=1}^{N_T} x_i^P,$$

where

$$\eta(\mathbf{x}^P) = \tilde{c} \gamma^{-1} (\tilde{\alpha} + \gamma)^{-1} \left(\gamma + \tilde{\alpha} e^{T(\tilde{\alpha} + \gamma)} - \gamma e^{\tilde{\alpha} T} - \tilde{\alpha} e^{\tilde{\alpha} T} \right) e^{-T(\tilde{\alpha} + \gamma) - \tilde{x}\tilde{\psi}}.$$

In this formulation, the rational steward anticipates the impact of its impact on the market for attacks as it sets x_i^P , rather than in the case of individual firms, where x_i^P is exogenous and set in equilibrium.

As mentioned above the possible differences between the discount rates of the steward and the individual agents are important factors in determining the optimal level of investment in information security. In the first instance, we evaluate the case where $\beta_i = \beta$ and $\delta = \beta$ which permits an exact identification of the Pareto optimal allocation of defensive expenditure and illustrates the total reduction in the present value of expected losses gained from achieving the Pareto optimal allocation versus the Nash equilibrium. Subsequently we examine the case where the steward sets a specific discount rate $\delta < \beta$. In this instance the steward values future security outcomes more highly than an individual firm.

With N_T ex-ante identical targets this reduces the policy-maker's problem to a representative one dimensional optimization problem. In this case, the attackers identify $x_i^P = x^P$, α , ψ , and c precisely (as targets are identical). Their problem does not change substantially, however, and the optimal attacker per target is defined in equilibrium as

$$\eta^P = \frac{\left(\gamma + \alpha e^{T(\alpha + \gamma)} - \gamma e^{\alpha T} - \alpha e^{\alpha T} \right) e^{-T(\alpha + \gamma) - x^P \psi}}{c \gamma (\alpha + \gamma)}.$$

The policy-maker takes fully into account the attacker intensity into their optimization problem. The algebraic expression is more complex, but the derivative with respect to x^P is analytic. The loss per target $L^P = TL^P / N_T$ over the time horizon $[0, T]$ is evaluated as

$$L^P = L(\alpha + \gamma)e^{T(\alpha + \gamma)} \frac{\left(\frac{Te^{-T(\alpha + \gamma) - \psi} \left(\alpha c \gamma \left(\gamma + \alpha e^{T(\alpha + \gamma)} - (\alpha + \gamma)e^{\alpha T} \right) + \delta(\alpha + \gamma)e^{T(\alpha + \gamma) + \psi} \right)}{\alpha + \gamma} \right)}{1 - e} \frac{1}{\delta(\alpha + \gamma)e^{T(\alpha + \gamma) + \psi} - \alpha c \gamma \left(-\gamma + \alpha \left(-e^{T(\alpha + \gamma)} \right) + (\alpha + \gamma)e^{\alpha T} \right)} + x.$$

This is the discounted present value using the policy-maker's discount rate and in the first instance $\delta = \beta$. Differentiating with respect to x yields the optimal level of investment for each firm under the guidance of the steward.

Proposition 2. For a steward setting mandatory defensive expenditure x^P for N_T ex-ante identical targets with discount rate δ on future expected losses, x^P is the solution for x of the following equation:

$$\begin{aligned} & \delta L \psi (\alpha + \gamma)^2 e^{2T(\alpha + \gamma) + \psi} \frac{\left(\frac{Te^{-T(\alpha + \gamma) - \psi} \left(\alpha c \gamma \left(\gamma + \alpha e^{T(\alpha + \gamma)} - (\alpha + \gamma)e^{\alpha T} \right) + \delta(\alpha + \gamma)e^{T(\alpha + \gamma) + \psi} \right)}{\alpha + \gamma} \right)}{\left(\alpha c \gamma \left(-\gamma + \alpha \left(-e^{T(\alpha + \gamma)} \right) + (\alpha + \gamma)e^{\alpha T} \right) - \delta(\alpha + \gamma)e^{T(\alpha + \gamma) + \psi} \right)^2} \\ & + \alpha c \gamma L T \psi \left(-\gamma + \alpha \left(-e^{T(\alpha + \gamma)} \right) + (\alpha + \gamma)e^{\alpha T} \right) \\ & \frac{e^{-\frac{\alpha c \gamma T \left(\gamma + \alpha e^{T(\alpha + \gamma)} - (\alpha + \gamma)e^{\alpha T} \right) e^{-T(\alpha + \gamma) - \psi}}{\alpha + \gamma}} - \delta T - \psi}{e} \\ & \frac{1}{\delta(\alpha + \gamma)e^{T(\alpha + \gamma) + \psi} - \alpha c \gamma \left(-\gamma + \alpha \left(-e^{T(\alpha + \gamma)} \right) + (\alpha + \gamma)e^{\alpha T} \right)} = 1. \end{aligned}$$

Although this equation is not analytically solvable for x^P , it is computable once the policy-maker and attackers' discount rates δ , γ are chosen along with the nominal loss L . The remaining terms are the technology parameters L and ψ and are subject to uncertainty.

We now set up the following simulations: the discount rate of the attackers $\gamma = 15\%$, the discount rate for both the steward and the targets $\delta = \beta = 25\%$; the ranges of the technology parameters are given as $\psi = \{0.1, 0.25, 0.5\}$ (defence is not effective to very effective), $\alpha \in \{0, 2\}$ (attacking effort is not effective to very effective), and $L (\approx \$500,000)$: on the basis of these we compute x^N versus X^P over the range of α . We then use these values of x^N and X^P to compute the nominal loss factor

$$R(x, \eta) = \int_0^T (1 - e^{-\alpha \eta t}) e^{-\psi x} dt$$

and the discounted expected total loss

$$TL = \int_0^T e^{-\delta t} L (1 - e^{-\alpha \eta t}) e^{-\psi x} dt + x,$$

for comparison purposes. The results are depicted in **Figure 4**.

We now consider the case where $\delta < \beta$ (see **Figure 5**); that is, the policy-maker has longer term time preferences than the ex-ante identical targets. It is here that the steward deviates from standard notions of a benevolent public policy-maker and this relates explicitly to the sustainability concept outlined above.

The steward's time preferences indicate longer horizon planning than the individual participants, by setting $T = -\log(1 - \lambda) \delta^{-1}$, for a value of λ close to one.

Using the previous case parameter values as a starting point we solve for x^P versus x^S whilst varying $\alpha \in \{0, 2\}$ and for $\psi = \{0.1, 0.25, 0.5\}$, for a group of N_T ex-ante identical targets.

The results are as expected: for all configurations of α and ψ considered, the level of defensive expenditure is higher. The nominal loss factor $R(x, \eta)$, the truly fair comparison between the Pareto steward versus the long term steward, indicates that total nominal expected loss per pound from attacking effort drops substantially as the steward sets a longer term strategy.

The ex-ante identical individual firms now no longer view the imposed x^S as Pareto optimal and losses included and balanced off by the steward are not viewed in the same way by the individual firm. Therefore, for

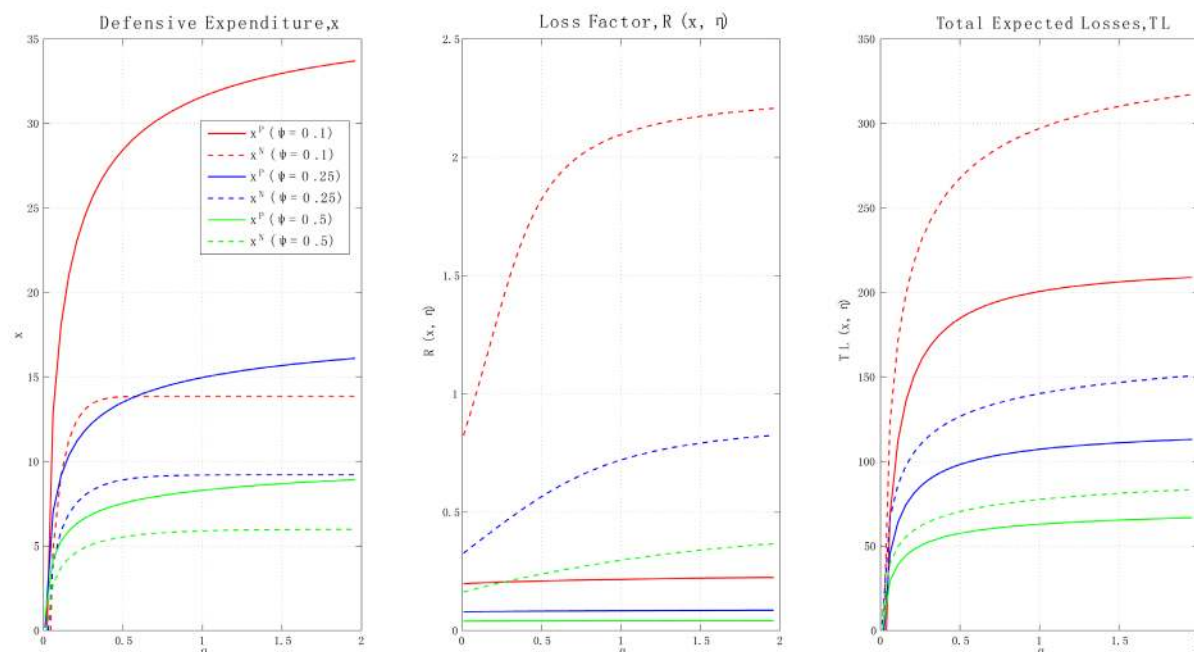


Figure 4. Comparison of the impact of a steward on defensive expenditure, risk, and expected losses assuming that the ecosystem consists of UK SMEs of the type surveyed in Case 2. For tractability, we assume identical targets. The level of defensive expenditure for each target is denoted by x (the left plot), the expected loss factor $R(x, \eta) = \int_0^T (1 - e^{-\alpha t}) e^{-\psi x} dt$ (centre plot) and total expected losses $TL = \int_0^T e^{-\delta t} L(1 - e^{-\alpha t}) e^{-\psi x} dt + x$ (the right plot). The dashed lines are the values of defensive expenditure, loss factor, and total expected losses over a varying security decay factor α (the abscissa values) in the absence of the steward. The red line presents the scenario when the security effectiveness factor ψ is equal to 0.1 (low effectiveness), the blue and green lines present the cases for $\psi = 0.25$ (intermediate effectiveness) and $\psi = 0.5$ (high effectiveness). The unbroken lines represent the same cases, ceteris paribus; however, there is now a steward coordinating defensive expenditure. The time preferences of the agents are as follows: we assume identical targets each with a discount rate of $\beta = 0.25$, a steward with discount rate $\delta = \beta = 0.25$ and attackers with a discount rate of $\gamma = 0.15$. The value of information security assets at risk at time t is assumed to be $\int_0^T e^{-\beta t} L dt = \$2 \text{ Million to } \pounds 555,555 \approx \$500,000$ and TL is presented per \$1000 of assets. Attacker rewards are set such that they receive 0.15 units of revenue per unit of effort therefore $c = 0.85$.

all values of α and ψ , the long-term steward is deemed to be more expensive in present value terms. In all cases, the size of this effect diminishes with decreasing α (reduced attacker efficacy) and increasing ψ (increasing defensive effectiveness). There is an externality created by the attacker-choice η which interacts with the new choice of δ . The steward views the externality as being larger compared to its evaluation by individual targets (as the attacker-choice of η lasts for the whole time horizon $t_0 < t < T$, the action of reducing δ necessarily increases T and subsequently the valuation of the externality).

Unambiguously, if the steward has a lower discount rate than the individual target firms, then the value of the costly action deemed necessary to negate the externality will be higher than that required by the targets to attain the Pareto optimal allocation subject to their time preferences.

A more attractive way of thinking about discount rates is to derive the time horizon over which the majority of their value amortizes towards zero. The firm-specific discount rates are set to amortize the firm's current investment assets over a time period consistent with the lifespan of previous information security assets. This provides a baseline for the steward's time horizon in terms of managing externalities. Should the steward desire the externalities to be managed over a longer, more sustainable, time horizon, then the adopted discount rate will be set lower than the representative rate determined by the individual firms. Larger scale ecosystems, such as the internet, are usually assumed to require longer term planning. Hence, stewards in this context might amortize expected losses from risks to the system over much longer periods. Therefore the incremental investment in

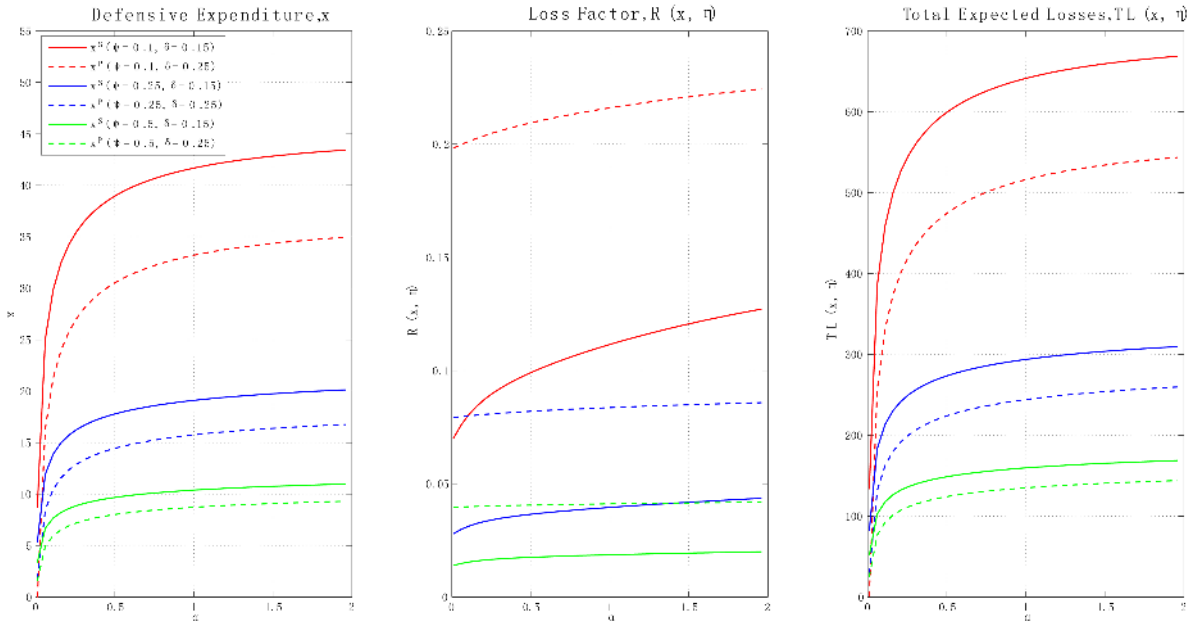


Figure 5. We compare the allocation of optimal investment x^P versus x^S , for the cases in which $\delta = \beta$ and $\delta < \beta$, respectively. The left hand plot presents the optimal allocation of defensive expenditure when $\delta = 0.25$ (dashed line) versus $\delta = 0.15$ (unbroken line) for varying security effectiveness ψ (red $\psi = 0.1$, blue $\psi = 0.25$ and green $\psi = 0.5$) against the security decay factor (the abscissa values) α varying from 0 (no decay) to 2 (attackers erode security very quickly). The centre plot presents the loss factor $R(x, \eta)$ and the right hand plot the total expected losses to targets assuming their own discount rate $\beta = 0.25$. Attackers have a discount rate of $\gamma = 0.15$. Information security assets at risk at time t is assume to be $\int_0^T e^{-\beta t} L dt = \$2 \text{ Million} \approx 500,000$ and TL is presented per \$1000 of assets. Attacker rewards are set such that they receive 0.15 units of revenue per unit of effort; therefore, $c = 0.85$.

information security that are imposed on individual participants in the ecosystem are deemed to be unnecessarily burdensome given their own time preferences.

If $\delta = \beta$ (see Figure 5), the resulting universal increase in defensive expenditure may not be sufficient to meet the sustainability target. In this case, the steward sets a discount rate less than β in order to achieve the sustainability target within the predefined time period. Such divergence of discount rates between the individual and the social is common in finance and environmental economics. An example of the debate on choice and imposition of social discount rates is in the climate change literature [15], where the choice of discount rate is particularly acute as the forward horizons are over multiple decades and centuries and, in this context, exponential discounting reduces future losses toward zero after a finite number of years. In the information security context, the impact of the time preference assumption is not so acute as investment horizons are much shorter (see for instance [16] for a model of investment horizons). However, the interaction of the externality with the differentiated discount rates between targets and the steward does indicate that this is an important issue for information ecosystems.

3. A Model of Resilience

A key feature of our account of resilience is that we illustrate how thresholds for the effectiveness of stewards emerge from the underlying model of the response of information ecosystems to the hostility of the environment.²

Again, we consider a set of N_T *ex-ante* identical targets choosing to allocate defensive resources that mitigate the harm from attacks. However, in this case, the targets need to solve, simultaneously, a multi-dimensional

²Note by that the hostility of the environment we mean a representation of the capacity of attackers rather than simply the success or failure of an individual attack.

resource-allocation problem. Let the subscripts h and l represent to potential areas of allocation of assets, where h and l denote the areas of high and low security at which information assets are held, and let $x_h \geq 0$ and $x_l \geq 0$ denote the one-off investments made at time t_0 in securing assets located in the corresponding areas. Finally, we define z to be a switching variable such that a fraction, $0 \leq z \leq 1$, of assets is allocated between h and l .

Our model depends crucially on two key (vector) parameters. Employing the same notation albeit in two dimensions h and l , we consider the elasticity of attacking intensity denoted by the vector α . This is the parameter that captures the marginal effectiveness of an additional attacker per target (η) entering the ecosystem.

In this case, we need to consider just two elasticities, α_l and α_h with corresponding η_l and η_h , which are associated with low and high levels of difficulty in securing assets. Second, we consider parameters ψ_l and ψ_h , which capture the relative rate of risk reduction for additional security investments by targets in each asset class (x_l and x_h).

Let $\tilde{\sigma}_{i \in \{l, h\}} : \mathbb{R}_+ \rightarrow [0, 1]$ be a function that determines the instantaneous time t risk for a fixed time-horizon, where $(t_0, T) = \{t | t_0 < t < T\}$. When properly specified we can interpret $\tilde{\sigma}$ as the instantaneous probability of a successful attack. We refer to z as the ‘‘asset allocation’’ and the quantities x_l and x_h as the ‘‘investment allocation’’, stated combinations of all three are referred as ‘‘allocation bundles’’.

Our assumption is that increased investment $x_{i \in \{l, h\}}$ reduces the probability of a successful attack; that is, $\partial \tilde{\sigma}_{i \in \{l, h\}} / \partial x_{i \in \{l, h\}} < 0$, ceteris paribus. However, along with increasing investment there is a decreasing marginal reduction in the probability of a successful attack, $\partial^2 \tilde{\sigma}_{i \in \{l, h\}} / \partial x_{i \in \{l, h\}}^2 > 0$. Similarly, with increased attacking intensity $\eta_{i \in \{l, h\}}$ on the particular area of allocation there should be a corresponding increase in the probability of a successful attack, $\partial \tilde{\sigma}_{i \in \{l, h\}} / \partial \eta_{i \in \{l, h\}} > 0$.

A functional form for $\tilde{\sigma}$ that satisfies these conditions is the following multiplicative model:

$$\tilde{\sigma}_i = e^{-\psi_i x_i} \eta_i^{\alpha_i}, i \in \{l, h\}. \quad (1)$$

Under this formulation, there is an upper bound on $\eta_{i \in \{l, h\}}$ of $\eta_i^* < e^{\alpha_i^{-1} x_i \psi_i}$, for $i \in \{l, h\}$, such that $\tilde{\sigma}_i$ may still be interpreted as probability of a successful attack. Here $\psi_{i \in \{l, h\}}$ is the relative marginal decrease in $\tilde{\sigma}_{i \in \{l, h\}}$ for a unit increase in $x_{i \in \{l, h\}}$ whilst $\alpha_{i \in \{l, h\}}$ is the elasticity of attack.

Let the number of attackers for each asset area be $N_{A, i \in \{l, h\}}$. The ratio of attackers per target is the attacking intensity $\eta_{i \in \{l, h\}} = N_{A, i \in \{l, h\}} / N_T$. Let the reward $R > 0$ for a successful attack be proportional to the assets allocated in each area, h and l , and for notational simplicity let $\zeta_{i=l} = z$ and $\zeta_{i=h} = 1 - z$. Set $\tilde{c} = c/R$ to be the cost ratio of attack, where c is the unit cost of a single attack. When the attacker’s time preference is described by γ , the profit function for a single attacker is

$$\tilde{\Pi}_{A, i} = \int_{t_0}^T e^{-\gamma t} \zeta_i \eta_i^{-1} \tilde{\sigma}_i(x_i, \eta_i) dt - \tilde{c}, i \in \{l, h\}. \quad (2)$$

We assume that attackers do not coordinate attacks (or are commissioned by a single attacker) and rewards are claimed on a first-winner-takes-all basis. Attackers are assumed to be drawn from a pool and make one-off entry decisions until marginal cost and marginal benefit are equal and hence $\tilde{\Pi}_{A, i} = 0$.

For the targets of such attacks, let $L > 0$ be an instantaneous value of assets at risk from attack and $\beta \in \mathbb{R}$ be a subjective discount rate determining the time preferences of all targets. The risk neutral expected loss over the time horizon $t_0 < t < T$, is given by

$$\tilde{V}_L = \int_{t_0}^T e^{-\beta t} (z \tilde{\sigma}_l(x_l, \eta_l) L + (1 - z) \tilde{\sigma}_h(x_h, \eta_h) L) dt + x_l + x_h. \quad (3)$$

The optimal allocation bundle $(z^\diamond, x_l^\diamond, x_h^\diamond)$, when attacking intensity is exogenous, is the simultaneous solution of $\{\partial \tilde{V}_L / \partial x_l = 0, \partial \tilde{V}_L / \partial x_h = 0, \partial \tilde{V}_L / \partial z = 0\}$. By construction, if $\alpha_{i \in \{l, h\}} > 0$, $\psi_{i \in \{l, h\}} > 0$, $L > 0$, $\beta > 0$ and $z \in (0, 1)$, a minimum of this function exists. By assumption we set that the optimal allocation must be either $(x_{i \in \{k, h\}}) \in \mathbb{R}_+$ when $(\eta_{i \in \{k, h\}}) \in \mathbb{R}_+$ or, if the minimum lies at $x_{i \in \{l, h\}} < 0$, then $x_{i \in \{l, h\}}^\diamond = 0$. Similarly, we impose the inequality constraint that $0 \leq z^\diamond \leq 1$.

In the case of resilience, the time horizon, T , is of greater significance compared to the previous case and maybe determined endogenously. Let λ be an arbitrarily large, but not infinite, number. For a given discount

rate, $\tilde{\theta} = \min(\delta, \beta)$, by construction:

$$\lim_{T \rightarrow \infty} \int_{t_0}^T \tilde{\theta}^{-1} e^{-\theta t} dt = 1.$$

Therefore, the approximation of the time horizon \tilde{T} covering the $1-1/\lambda$ proportion of the future losses is derived from $\tilde{T} = \log(\lambda)/\tilde{\theta}$. In this paper, we follow [17] and assume that $\beta > \gamma$ and $\tilde{T} = \log(\lambda)/\delta$, such that the interval t_0 to \tilde{T} covers 90% of the expected present value; that is, $\lambda = 10$.

What is important, to the steward, is the overall mass of attacks against systems containing assets under the types l and h “storage/operations” areas and this will be influenced by the aggregate behaviour of targets and attackers, rather than the microstructure of individual attack-defence interactions. The more attractive the ecosystem is to attackers, the greater the mass of attacks against its individual components.

Proposition 3. (Existence of Nash equilibria without the steward)

1) (Equilibrium Target Investment) Under the preceding assumptions, when $\tilde{\sigma}_i = e^{-\psi_i x_i} \eta_i^{\alpha_i}$, for $i \in \{l, h\}$, the Nash equilibrium allocations of x_h , x_l and z denoted x_h^* , x_l^* and z^* are

$$x_i^* = \frac{\alpha_i}{\psi_i} \log \left(\frac{L \psi_i \psi_j^2 (e^{\delta T} - 1)^2}{\tilde{c} \delta \beta (\psi_j + \psi_i)^2} \right) - \frac{\alpha_i \delta T}{\psi_i}, \quad i \in \{l, h\}, j \in \{l, h\}, j \neq i$$

$$z^* = \frac{\psi_l}{\psi_h + \psi_l}. \quad (4)$$

2) (Equilibrium Attacker Intensity) Following from Part 1, above, the Nash equilibrium attacker intensities, denoted η_l^* and η_h^* are

$$\eta_i^* = \left(\frac{\psi_j (e^{\delta T} - 1) e^{-x_i^* \psi_i - \delta T}}{\tilde{c} \delta (\psi_i + \psi_j)} \right)^{\frac{1}{1-\alpha_i}}, \quad i \in \{l, h\}, j \in \{l, h\}, j \neq i, \quad (5)$$

where $x_{i \in \{l, h\}}^*$, is the functional forms of the Nash equilibrium given in Part 1 (above).

We demonstrate that, in this modelling approach, we do not have to impose an arbitrary constraint on $x_l + x_h$, to create conditions similar to the standard results obtained when optimizing under such budget restrictions.

3.1. Introducing the Steward

The aim of the steward here is to ensure resilience, as defined above: the information system may not, given the choices of investment in information security allocated to the individual components, be resilient to debilitating shocks.

The first stewardship action that we evaluate replicates our previous work (on sustainability [17]) by postulating a Stackelberg policy framework in which the policy-maker stewarding the system sets rules relative to a target level of sustainability. When the steward is fully informed, our model reverts to the mechanism design problem discussed in [17], in which the steward is able to set a mandatory investment bundle (denoted by the lower bar) on the individual targets (\bar{x}_l, \bar{x}_h) , as well as imposing a specific asset allocation \bar{z} .

The Nash equilibrium allocation for the N_T targets assumes no social coordination. Therefore, the Nash equilibrium allocation (x_l^*, x_h^*, z^*) of defensive effort and corresponding attacking intensities, (η_l^*, η_h^*) , will not necessarily be the best solution for Pareto efficiency. Let $(x_l^\dagger, x_h^\dagger, z^\dagger)$ be the Pareto efficient allocations for a given set of model parameters $(\alpha_{i \in \{l, h\}}, \beta, \tilde{c}, \delta, \lambda, \psi_{i \in \{l, h\}}, L)$.

A classical efficiency model, with the steward acting as a public policy-maker and imposing $(\bar{x}_l, \bar{x}_h, \bar{z})$ [17], demonstrates that Pareto efficiency is guaranteed only when the subjective discount rate of the steward is equal to β (the common discount rate).

Indeed, the analysis of sustainability [1] [2] illustrates that, from the subjective viewpoint described by targets' heterogeneous discount rates, the chosen values of $(\bar{x}_l, \bar{x}_h, \bar{z})$ cannot always be a Pareto efficient allocation, $(x_l^\dagger, x_h^\dagger, z^\dagger)$, when $\beta \neq \bar{\beta}$. However, there may exist constellations of parameters such the welfare of the individual agents have improved due to the presence of the steward despite their different discount rates.

Let the steward's discount rate be $\bar{\beta}$. A fully informed steward sets a mandatory level of $(\bar{x}_l, \bar{x}_h, \bar{z})$ by

minimizing the following loss function

$$\tilde{V}_p = \int_{t_0}^T e^{-\bar{\beta}t} \left(z \tilde{\sigma}_l(x_l, \eta_l^*) L + (1-z) \tilde{\sigma}_h(x_h, \eta_h^*) L \right) dt + x_l + x_h, \quad (6)$$

where $\eta_i^*(x_i, z)$ for $i \in \{l, h\}$ is the solution to

$$\int_{t_0}^T e^{-\delta t} \zeta_i \eta_i^{-1} \tilde{\sigma}_i(x_i, \eta_i) dt = \gamma, \quad i \in \{l, h\}, \quad (7)$$

in terms of (x_l, x_h, z) . As in the previous case by internalizing the attacker reaction curve, the fully informed policy-maker with identical time preferences to the homogenous targets $\bar{\beta} = \beta$ will set an allocation bundle $(\bar{x}_l, \bar{x}_h, \bar{z})$.

Proposition 4. (The fully informed steward)

1) (Target investment with steward) When $\tilde{\sigma}_i = e^{-\psi_i x_i} \eta_i^{\alpha_i}$ and $\bar{z} = \psi_h / (\psi_h + \psi_l)$, the stewards optimal investment allocation (\bar{x}_l, \bar{x}_h) is

$$\begin{aligned} \bar{x}_i &= \frac{1}{\psi_i} \log \left(\psi_j (\psi_i + \psi_j)^{\frac{1}{1-\alpha_j}} \right) + \frac{\alpha_i}{\psi_i} \log \left(\frac{1}{\gamma} \delta (e^{\delta T} - 1) \right) \\ &+ \left(\frac{\bar{\beta} T (\alpha_i - 1)}{\psi_i} - \frac{\delta T \alpha_i}{\psi_i} \right) + \frac{(\alpha_i - 1)}{\psi_i} \log \left(\frac{-\bar{\beta} (\alpha_j - 1)}{L \psi_i (e^{\bar{\beta} T} - 1)} \right) \end{aligned} \quad (8)$$

$i \in \{l, h\}, j \in \{l, h\}, j \neq i.$

2) (Attacking intensity) Following from Part 1, above, the attacker intensity $\eta_{i \in \{l, h\}}$ is

$$\bar{\eta}_i = \left(\frac{\psi_i (e^{\delta T} - 1) e^{-\bar{x}_i \psi_i - \delta T}}{\gamma \delta (\psi_j + \psi_i)} \right)^{\frac{1}{1-\alpha_i}} \quad i \in \{l, h\}, j \in \{l, h\}, j \neq i, \quad (9)$$

where \bar{x}_i is given in Equation (8).

Proposition 5. (The steward's improvement) If $\tilde{\sigma}_i = e^{-\psi_i x_i} \eta_i^{\alpha_i}$, with $\beta \geq \bar{\beta}$, and $\alpha_{i \in \{l, h\}} > 0, \psi_{i \in \{l, h\}} > 0, \gamma > 0, \delta > 0, L > 0$ and the asset allocation is constrained to $\bar{z} = \psi_h / (\psi_h + \psi_l)$, then the steward's mandated investment $\bar{x}_{i \in \{l, h\}}$ is always greater than or equal to the Nash equilibrium investment bundle $x_{i \in \{l, h\}}^*$.

A useful by-product of the comparison between Propositions 1 and 2 is that we can define an upper bound on $\beta \geq \bar{\beta}$ such that the steward does at least as well as the Nash equilibrium even when the steward weights potential near-term losses more than the targets do. Again, this is covered in more detail for the one-dimensional case in [17].

We now progress to the case of a partially informed steward with a minority action: Let the steward observe and enforce only x_h . The steward can observe and internalize the externality in η_h , but cannot observe or enforce z or x_l . The targets then choose the investment and allocation bundle (x_l, z) following:

$$(\tilde{x}_l, \tilde{z}; x_h, \eta_l, \eta_h) = \underset{x_l, z}{\operatorname{argmin}} \int_{t_0}^T e^{-\beta t} \left(\bar{z} \tilde{\sigma}_l(x_l, \eta_l) L + (1-\bar{z}) \tilde{\sigma}_h(x_h, \eta_h) L \right) dt + x_l + x_h \quad (10)$$

and the steward now solves the following minority optimization, with the steward's given information set:

$$\bar{x}_h(\eta_h) = \underset{x_h}{\operatorname{argmin}} \int_{t_0}^T e^{-\beta t} \left(\hat{L} \tilde{\sigma}_h(\bar{x}_h, \eta_h^*) \right) dt + x_h, \quad (11)$$

where η_h^* is the solution to the attacker entry problem from Equation (2), but only for the h asset class.

From the steward's point of view, this is now

$$\int_{t_0}^T e^{-\delta t} \tilde{\zeta}_h \eta_h^{-1} \tilde{\sigma}_h(x_h, \eta_h) dt = \tilde{c}. \quad (12)$$

Note that the steward now takes for given L as the value of losses. This is because the steward can no longer identify zL and $(1-z)L$; the steward is simply given \hat{L} by the targets a priori, and this is assumed to be exogenous. Similarly, whilst $\tilde{\zeta}_h$ is equal to z from the viewpoint of attackers and targets, it is simply a parameter unrelated to the overall asset allocation of the targets from the point of view of the steward. The steward is now, unwittingly, not a Stackelberg policy-maker, but in a Nash equilibrium with the targets and attackers.

Proposition 6. (Attackers and the steward)

1) (Asset Class h investment guided by the steward) If $\tilde{\sigma}_i = e^{-\psi_i x_i} \eta_i^{\alpha_i}$, for $i \in \{l, h\}$, the steward's objective function is as stated in Equation (11), and the attacker dynamics are as given in Equation (12), then the steward's optimal mandated investment allocation is

$$\begin{aligned} \bar{x}_h = & \frac{1-\alpha_h}{\psi_h} \log(\hat{L}(e^{\beta T} - 1)\psi_h) - \frac{\alpha_h}{\psi_h} \log(\gamma\delta(\tilde{\zeta}e^{\delta T} - \tilde{\zeta})) \\ & + \frac{1}{\psi_h} (\log(\bar{\beta}\alpha_h - \bar{\beta})(1-\alpha_h) + \alpha_h T(\bar{\beta} - \delta) - \bar{\beta}T). \end{aligned} \quad (13)$$

Following from the steward's choice, the attacker intensity, given the steward's actions $\bar{\eta}_h$, is given by

$$\bar{\eta}_h = \left(\frac{\tilde{\zeta}(e^{\delta T} - 1)e^{-\bar{x}_h\psi_h - \delta T}}{\gamma\delta} \right)^{\frac{1}{1-\alpha_h}}, \quad (14)$$

where \bar{x}_h is as defined in Equation (13). Investment in class l will be decided by the agents, conditional on the steward's recommendation for h .

2) (Asset Class l) We now consider the targets' and attackers' new equilibrium: if $\tilde{\sigma}_i = e^{-\psi_i x_i} \eta_i^{\alpha_i}$, for $i \in \{l, h\}$, and the targets' objective is as specified in Equation (10), then the equilibrium allocation bundle x_l, z will be

$$\begin{aligned} x_l^\ddagger = & -\frac{1}{\psi_l} \log(\bar{\eta}_h^{\alpha_h}) \\ & + \frac{\alpha_l}{\psi_l} \left(\log(\bar{\eta}_h^{\alpha_h}) + \log(\beta(e^{\delta T} - 1)\bar{\eta}_h^{-\alpha_h}) - \log(\gamma\delta\psi_l(e^{\beta T} - 1)L) + \beta T - \delta T + \bar{x}_h\psi_h \right) \\ z^\ddagger = & \frac{\beta\bar{\eta}_h^{-\alpha_h} e^{\bar{x}_h\psi_h + \beta T}}{L\psi_l(e^{\beta T} - 1)}, \end{aligned} \quad (16)$$

and the attacker intensity η_l is given by

$$\eta_l^\ddagger = \left(\frac{z(e^{\delta T} - 1)e^{-x_l^\ddagger\psi_l - \delta T}}{\gamma\delta} \right)^{\frac{1}{1-\alpha_l}}. \quad (17)$$

Note that we use the \ddagger to denote this new equilibrium for the targets as it is not strictly a Nash equilibrium solution, but rather is Bayes-Nash equilibrium, in which the steward has prior values for \hat{L} and $\tilde{\zeta}$. See [18] for an explanation of the difference between Nash and Bayes-Nash equilibria.

3.2. Measuring Resilience

Measuring the impact of technological shocks to $\alpha_{i,i \in \{l, h\}}$ and $\psi_{i,i \in \{l, h\}}$ and economic shocks to $\bar{\beta}$, β , δ , L , and \tilde{c} is a challenging task which requires the creation of an arbitrary metric. Here, we combine the equilibrium values of $x_{i \in \{l, h\}}$, z , and $\eta_{i \in \{l, h\}}$ using a total non-discounted loss function for the risk component only. This is given as follows:

$$\tilde{V}_A(\tilde{v}, \tilde{u}) = \int_{t_0}^{\tilde{T}} \tilde{z} \tilde{\sigma}_l(\tilde{x}_l, \tilde{\eta}_l) L + (1 - \tilde{z}) \tilde{\sigma}_h(\tilde{x}_h, \eta_h) L dt \quad (18)$$

$$\tilde{v} = \left(\tilde{z}, \tilde{x}_{i \in \{l, h\}}, \tilde{\eta}_{i \in \{l, h\}} \right) \quad (19)$$

$$\tilde{u} = \left(\alpha_{i \in \{l, h\}}, \psi_{i \in \{l, h\}} \right), \quad (20)$$

where $\tilde{T} = -\log(\lambda)/\theta$ and $\theta = \min(\bar{\beta}, \beta, \delta)$, for an arbitrary λ tending to zero. By construction, Equation 18 gives an undiscounted loss function, so that the value of the critical parameter \tilde{T} , which represents the step-size of the periods considered in the model (cf. **Figure 3**, for a multi-period sustainability model), is finite. \tilde{v} is the collection of choice variables under the various stewardship options and \tilde{u} is the collection of parameters that are subject to the technology shocks under consideration.

For a single period, resilience will be measured by a response function to shocks to the parameters \tilde{u} . Our choice of the response function for technology shocks allows for shocks across the set of parameters, \tilde{u} , either simultaneously or individually. It is given by the numerical evaluation of the following ordinary differential equation:

$$\tilde{I}(\tilde{u}) = \int_{t_0}^{\tilde{T}} \frac{\partial \tilde{z}}{\partial \tilde{u}} \tilde{\sigma}_l \left(\frac{\partial \tilde{x}_l}{\partial \tilde{u}}, \frac{\partial \tilde{\eta}_l}{\partial \tilde{u}} \right) L + \frac{\partial(1-\tilde{z})}{\partial \tilde{u}} \tilde{\sigma}_h \left(\frac{\partial \tilde{x}_h}{\partial \tilde{u}}, \frac{\partial \tilde{\eta}_h}{\partial \tilde{u}} \right) L dt, \quad (21)$$

$$\tilde{u} = \left\{ \alpha_{i \in \{l, h\}}, \psi_{i \in \{l, h\}} \right\},$$

where each case has a set of functional forms for \tilde{z} , $\tilde{x}_{i \in \{l, h\}}$ and $\tilde{\eta}_{i \in \{l, h\}}$. We have denoted the three cases as follows: $\tilde{v} = \tilde{v}^*$ and $\tilde{u} = \tilde{u}^*$ for the Nash equilibrium, $\tilde{v} = \tilde{v}^{\dagger}$ and $\tilde{u} = \tilde{u}^{\dagger}$ for the fully informed steward and $\tilde{v} = \tilde{v}^{\ddagger}$ and $\tilde{u} = \tilde{u}^{\ddagger}$ for the partially informed steward with minority action case.

We are interested in establishing the thresholds, illustrated in **Figure 6**, which describe levels of system operating capacity, as measured by loss, for differing degrees of the steward's effectiveness. We attempt to establish whether the system restores, through co-ordinated investment, to the target zone or not.

In our model, these thresholds reveal themselves as discontinuities, relative to shock size, in the solutions to Equation (21), below. Such discontinuities can be seen in our simulations as the asymptotes in **Figure 7** and **Figure 8**, for fully and partially informed stewards.

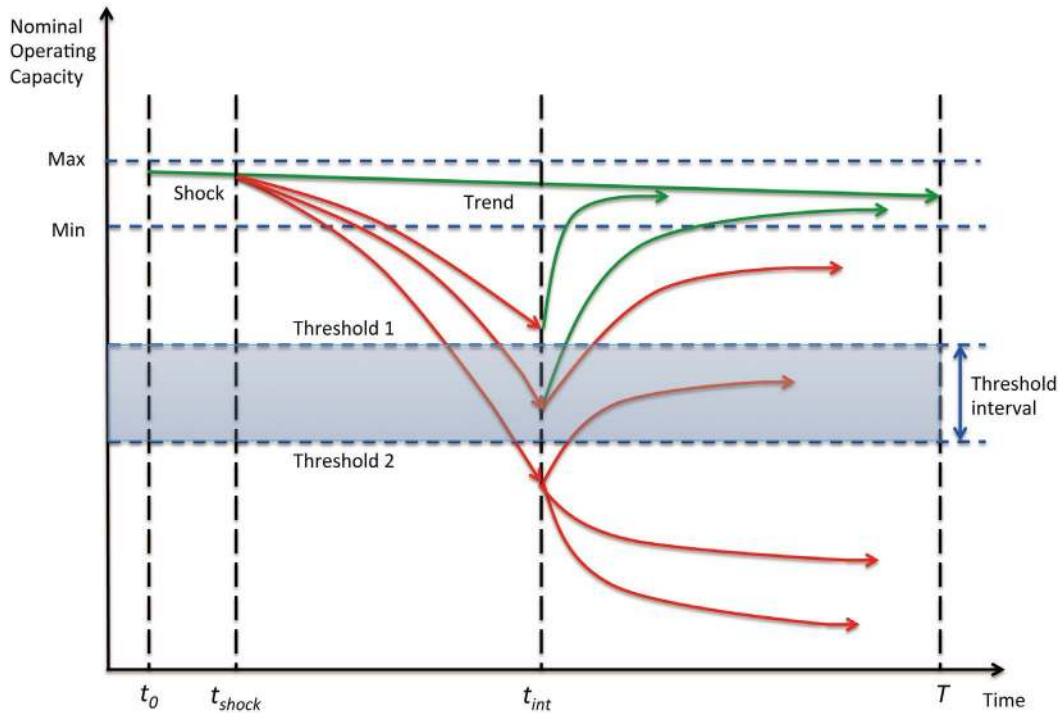


Figure 6. Resilience with an incompletely informed steward.

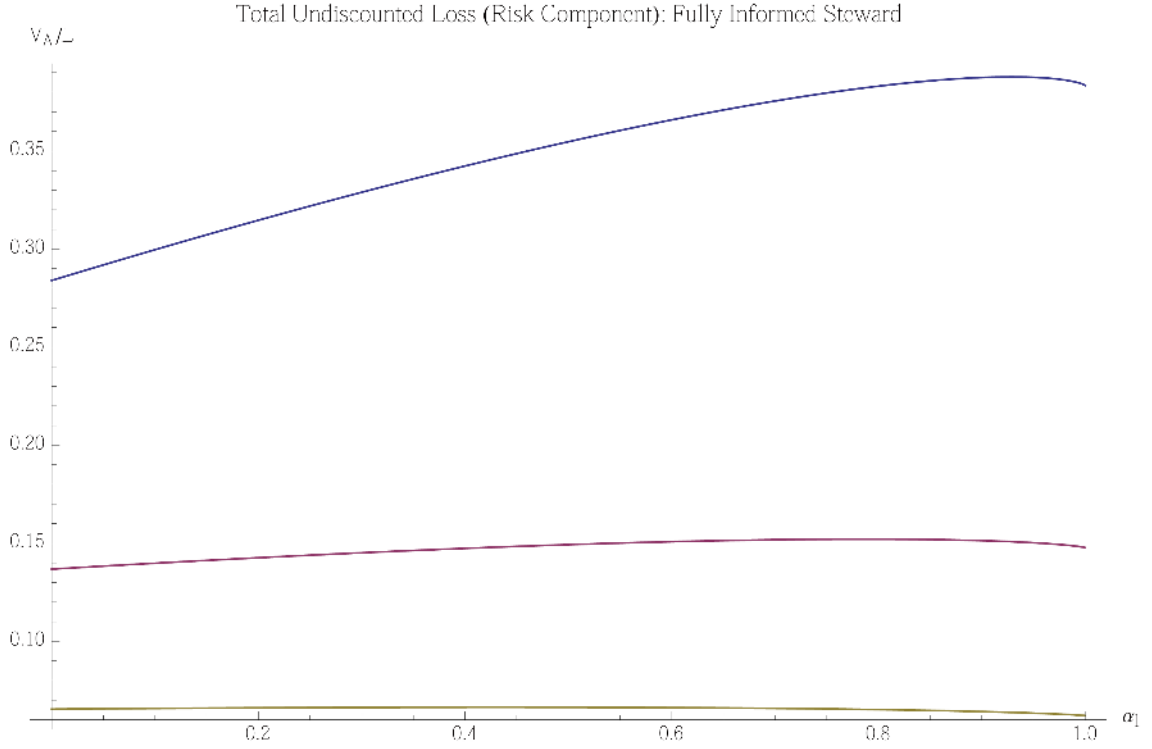


Figure 7. The steward's total non-discounted loss function, \tilde{V}_A , as a function of α_1 . An important point to note is that this does not include the deterministic up-front investment, so this curve can actually slope downwards, even with increasing α_1 . The upper curve represents $\psi_l = \psi_h = \psi = 0.01$, the middle curve $\psi = 0.1$ and the lower curve is $\psi = 0.5$. These values of ψ represent, respectively, low, medium, and high rates of risk reduction for additional investment.

Let $\tilde{V}_A(v^*, u^*)$ and $\tilde{I}(u^*)$ be, respectively, the total non-discounted loss for the risk component under the Nash equilibrium and the corresponding collection of response functions. Similarly let $\tilde{V}_A(\bar{v}, \bar{u})$, $\tilde{I}(\bar{u})$ and $\tilde{V}_A(\bar{v}^\ddagger, \bar{u}^\ddagger)$, $\tilde{I}(\bar{u}^\ddagger)$ be, respectively, the same pair of functions and collection of functions for the fully informed steward and the partially informed steward with minority action cases.

We can measure the effectiveness of the steward by comparing $\tilde{V}_A(v^*, u^*)$ to $\tilde{V}_A(\bar{v}, \bar{u})$. We can also evaluate the erosion in risk reduction caused by restricting the stewards information set and action space by pairwise evaluation of $\tilde{V}_A(v^*, u^*)$ and $\tilde{V}_A(\bar{v}, \bar{u})$ with $\tilde{V}_A(\bar{v}^\ddagger, \bar{u}^\ddagger)$.

To examine the impact of shocks and measure resilience we compare the response functions $\tilde{I}(u^*)$ and $I(\bar{u})$ to evaluate the impact of the fully informed steward. Finally, we can compare the resilience of the system when the stewards information set is restricted by comparing $\tilde{I}(u^*)$ and $I(\bar{u})$ to $I(\bar{u}^\ddagger)$, for varying sizes of shocks in \tilde{u} . In particular, we focus on $\alpha_{i \in \{l, h\}}$.

3.3. An Example Simulation

This simulation is designed to provide an overview of the intuition of our model and is not supposed to provide specific quantification for our proposed application. However, we have tried to stay close to real data when possible.

Let us assume that targets have a discount rate of 20% per annum ($\beta = \log(6/5)$ continuous growth rate), in this case when $\lambda = 10$, the target time overall horizon is $T = 12.3$ years. This appears to be a reasonable assumption for the amortization of information assets within a firm see, for example, the survey in [19]. For electricity transmission in the United States, the difference between physical and information assets can be found in [20] [21].

We assume that the societal discount rate used by the steward is much lower and ranges from $\bar{\beta} \rightarrow 0$ to

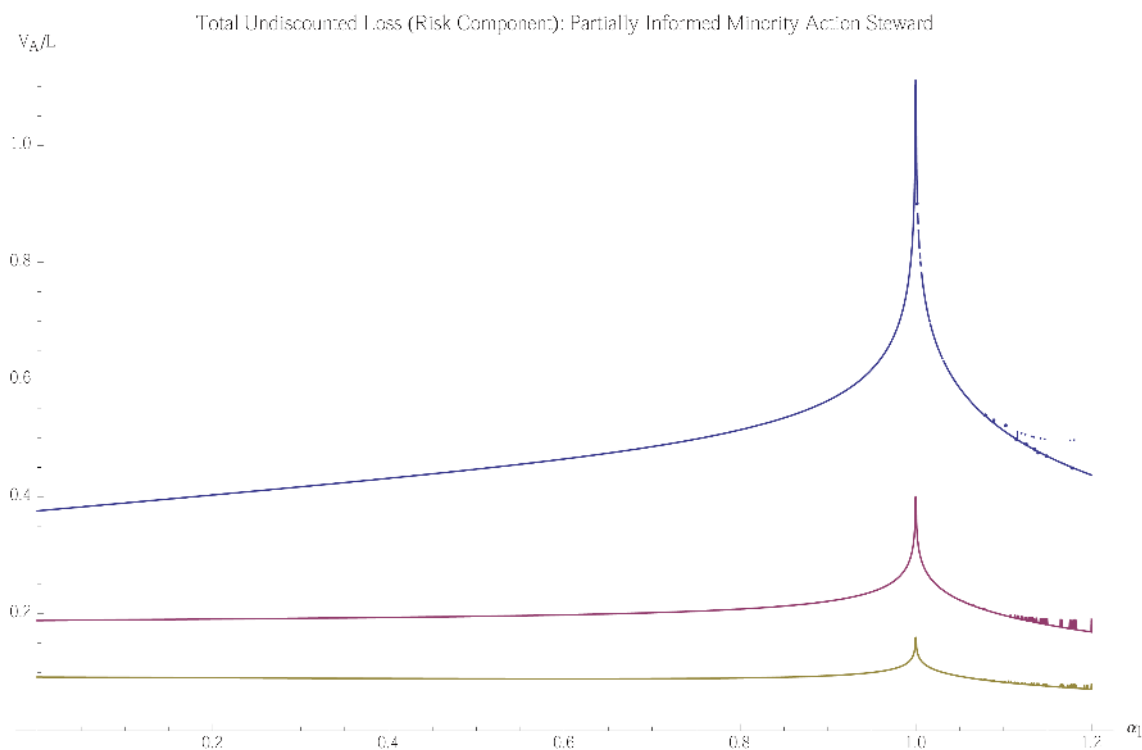


Figure 8. Partially informed steward with minority action total non-discounted loss function \tilde{V}_A as a function of α_l . In this case, the targets maintain assets in the increasingly risky l class to avoid costly regulation in h . However, a discontinuity exists at α_l , causing the loss function to spike before the assets are shifted back to the regulated domain.

$\bar{\beta} \rightarrow 1/10$. In [17], we outline the various debates on the appropriate social discount factor to be applied in public policy scenarios. For certain areas of public policy debate such as climate change discount rates approaching zero a common for certain economic arguments relating to low carbon policies. For information stewardship the requirement is not so acute but significant differences between firm discount rates and societal discount rates remain.

For our starting numerical example, we assume that $\psi_l = \psi_h$; that is, the relative marginal risk reduction from investment in both asset classes is identical and fixed. We assume that it is 1/100, 1/10 and 1/2, to represent low, medium and high effectiveness bands, respectively. This is a more difficult assumption to justify as there is very little literature on the efficacy of investment in security in this area; therefore, our simulation covers a wide range of reasonable bands.

We arbitrarily fix the instantaneous total loss to \$1M, as an example, and divide all losses by L to give a per-dollar-at-risk measure. Starting from the Nash equilibrium assumption, we presume that the rate of reduction in risk for a given investment is the same for both asset classes, it follows that the chosen weighting in each class will be 1/2 from the optimal weighting requirement.

We set the attackers' discount rate to be $\log(11/10)$, or a 10% discrete rate of return. From the viewpoint of attackers, the discount rate is analogous to an investment, as opposed to depreciation and amortization from the viewpoint of the targets. The most difficult parameter to set in the simulation is the cost per attack to reward ratio as almost no data exists on this quantity. When the cost to reward approaches zero, the cost per attack divided by reward indicates that either the rewards are very high or that the cost per attack is very low and intensity tends to infinity. This asymptotic trend in attacking has not been observed, therefore we stick to a ratio of about 10%. The shock of interest is that to the elasticity of attack $\alpha_{i \in \{l, h\}}$ and, in particular, shocks to α_l , hence attackers gain a new capability that allows them to be far more effective than before. This is the most realistic scenario—as opposed to significant shocks to the economy of attack or the capability of defence.

For the partially informed steward with minority action, the total non-discounted loss In this case, the pattern is similar to the Nash equilibrium for small shocks. The targets, however have costly regulation in the h asset

class and are under investing in the l asset class. Unfortunately, in this case there is a discontinuity at $\alpha_l = 1$, so the total loss spikes prior to the shift in assets from l to h . This is a de facto boundary, as illustrated in **Figure 6**. We can see that before the steward can regulate the assets, the total risk will traverse the discontinuity, before the steward can actually manage the majority of assets that the targets have not declared. Here, we can see a case of an ecosystem that is not resilient and lies within the feasible boundaries of our example parameter sets.

4. Conclusions

An information security ecosystem consists in a finite set of interacting agents supported by a specific infrastructure, which may have logical, physical, and economic components, supporting people, process, and technology.

The security posture of an ecosystem is a function of the postures of the participating agents and specified properties of the infrastructure that supports the ecosystem. The ecosystem exists and operates in an environment in which threats to the system's confidentiality, integrity and availability are initiated by a variety of sources. Such threats are taken as giving rise to shocks deteriorating ecosystem's operational posture. That is, its *sustainability*—the tendency of a firm or of the ecosystem to remain within acceptable operating bounds—and *resilience*—the tendency of an agent or of the ecosystem to restore, in a timely manner, its operating capacity to within acceptable operating range having been subject to a shock that has placed its operating capacity outwith it.

In this paper, we have addressed the question of whether it is socially desirable to introduce an institution with the authority to mandate resources for investment in information security in order to maintain (or preserve) the system's sustainability and resilience, rather than protecting these characteristics at the individual agents' levels. This institution, the *information steward* is therefore entrusted and empowered with the authority to maintain the sustainability and resilience of the security posture of the ecosystem. The steward's preferences are understood to derive in some form from the collective preferences of the participating agents, in their desire to maintain the system's operational stability.

The threats to the system's stability emanate from rational agents, the attackers, who, in deploying their resources, are fully informed of their potential gains given their technologies and, more importantly, the system's average protective posture. Agents operating in the system also invest in information security measures that take into account the costs of such investment and the possible losses following a successful attack. Both the defending agents and attackers engage in a situation of strategic interaction, and it is in this context that the relevant resources are deployed.

The introduction of the information steward, with authority to mandate to all agents in the system investment in information security, whilst fully informed on both the attacking and defending technologies, alters the decision landscape. In terms of maintaining and preserving the system's sustainability in the presence of secular deterioration, we show that the presence of the steward will act as a deterrent to attackers who will now know that the system's average defensive expenditure is higher thus reducing their expected gains from a successful attack by reducing markedly the probability of such event. In this case, the intensity of attacks declines and the system's secular decline is interrupted as its capacity shifts to higher level. The steward is shown to be effective in this respect. However, whether such intervention is strictly welfare-improving for all individual agents in the system is debatable, given the possibly large range of discount factors used by the agents in computing their "equilibrium" investment in information security. It is not uncommon to postulate that such a collective body will, in deciding the optimal investment in information security, use higher discount factors than the individual agents as its brief is the system's sustainability over some longer horizon, which may not coincide with the horizon of the agents, who aim at achieving short-run performance measures.

In terms of improving the system's ability to recover quickly after a negative temporary shock to its operational capacity—that is, of resilience—the influence of the presence of the information steward requires the appropriate design of the steward's mandate. There now two sets of strategic interactions: the usual one between the attackers and defenders, but, in addition, agents now have incentives to avoid additional information security investment—that which may be demanded by the steward—by deliberately mis-classifying the status of information assets which are under threat of attack. In such incomplete design, when the steward's ability to be fully informed is impaired and fails to enjoy the complete set of instruments at its disposal, the public co-ordination of investment in information security is actually welfare-reducing compared to the un-coordinated outcome for the

system as a whole.

In conclusion, public co-ordination of investment in information security in order to maintain the system's capacity in terms of sustainability and resilience can, by increasing the system's defensive posture, be effective in deterring attacks. It is important that a great deal of attention is required in designing such a co-ordinating institution, which will enjoy rather wide powers to decree the allocation of resources to information security. We have addressed the issue of the desirability of public co-ordination of investment in information security by bringing together a body of work on the common thread of the information steward. Integrating in a coherent manner these diverse results provides for a lucid understanding of the institutional requirements for the design and function of such a public institution and, depending upon the nature of the threat facing the information ecosystem, helps clarify the policy responses.

References

- [1] Ioannidis, C., Pym, D. and Williams, J. (2013) Sustainability in Information Stewardship: Time Preferences, Externalities, and Social Co-Ordination. In: Friedman, A., Ed., *Proceedings of the 12th Annual Workshop on the Economic of Information Security (WEIS 2013)*, Georgetown University, Washington DC, 11-12 June 2013. <http://weis2013.econinfosec.org/papers/IoannidisPymWilliamsWEIS2013.pdf>
- [2] Ioannidis, C., Pym, D. and Williams, J. (2014) Sustainability in Information Stewardship: Time Preferences, Externalities, and Social Co-Ordination. University College London, Department of Computer Science, Research Note RN/14/15. http://www.cs.ucl.ac.uk/fileadmin/UCL-CS/research/Research_Notes/rn-14-15_01.pdf
- [3] Ioannidis, C., Pym, D., Williams, J. and Gheyas, I. (2013) Resilience in Information Stewardship. In: Grossklags, J., Ed., *Proceedings of the 13th Annual Workshop on the Economic of Information Security (WEIS 2014)*, Pennsylvania State University, 23-24 June 2014. <http://weis2014.econinfosec.org/papers/Ioannidis-WEIS2014.pdf>
- [4] Ioannidis, C., Pym, D., Williams, J. and Gheyas, I. (2014) Resilience in Information Stewardship. University College London, Department of Computer Science, Research Note RN/14/16. http://www.cs.ucl.ac.uk/fileadmin/UCL-CS/research/Research_Notes/rn-14-16_01.pdf
- [5] Nardi, B. and O'Day, V. (1999) *Information Ecologies*. MIT Press.
- [6] Chapin III, F.S., Kofinas, G.P. and Folke, C. (2009) *Principles of Ecosystem Stewardship: Resilience-Based Natural Resource Management in a Changing World*. Springer-Verlag.
- [7] Stern, N. (2006) *Stern Review on the Economics of Climate Change: Executive Summary Long*. HM Treasury Stationary Office.
- [8] Hall, C., Anderson, R., Clayton, R., Ouzounis, E. and Trimintzios, P. (2013) Resilience of the Internet Interconnection Ecosystem. In: Schneier, B., Ed., *Economics of Information Security and Privacy III*, Springer, 119-148. http://dx.doi.org/10.1007/978-1-4614-1981-5_6
- [9] Benabou, R. and Tirole, J. (2012) *Laws and Norms*. Working Paper IZA DP No. 6290.
- [10] Funk, P. (2007) Is There an Expressive Function of Law? An Empirical Analysis of Voting Laws with Symbolic Fines. *American Economic Review*, **9**, 135-139. <http://dx.doi.org/10.1093/aler/ahm002>
- [11] Tyran, J. and Feld, L. (2006) Achieving Compliance When Legal Sanctions Are Non-Deterrent. *Scandinavian Journal of Economics*, **108**, 135-156. <http://dx.doi.org/10.1111/j.1467-9442.2006.00444.x>
- [12] Andreoni, J. (1989) Giving with Impure Altruism: Applications to Charity and Ricardian Equivalence. *Journal of Political Economy*, **97**, 1447-1458. <http://dx.doi.org/10.1086/261662>
- [13] Deci, E. (1985) *Intrinsic Motivation in Human Behavior*. Plenum. <http://dx.doi.org/10.1007/978-1-4899-2271-7>
- [14] Gordon, L. and Loeb, M. (2002) The Economics of Information Security Investment. *ACM Transactions on Information and Systems Security*, **5**, 438-457. <http://dx.doi.org/10.1145/581271.581274>
- [15] Caplin, A. and Leahy, J. (2004) The Social Discount Rate. *Journal of Political Economy*, **112**, 1257-1268. <http://dx.doi.org/10.1086/424740>
- [16] Ioannidis, C., Pym, D. and Williams, J. (2012) Fixed Costs, Investment Rigidities, and Risk Aversion in Information Security: A Utility-Theoretic Approach. In: Schneier, B., Ed., *Economics of Security and Privacy III*, Springer, Proceedings of the 2011 Workshop on the Economics of Information Security.
- [17] Ioannidis, C., Pym, D.J. and Williams, J.M. (2013) Sustainability in Information Stewardship: Time Preferences, Externalities, and Social Co-Ordination. *The Twelfth Workshop on the Economics of Information Security (WEIS 2013)*. <http://weis2013.econinfosec.org/papers/IoannidisPymWilliamsWEIS2013.pdf>
- [18] Fudenberg, D. and Tirole, J. (1991) *Game Theory*. MIT Press.

- [19] Baldwin, J., Gellatly, G., Tanguay, M. and Patry, A. (2005) Estimating Depreciation Rates for the Productivity Accounts. Technical Report, OECD Micro-Economics Analysis Division Publication.
- [20] Publications, N. (2013) Second Draft 2014 Business Plan and Budget. Technical Report, North American Electric Reliability Corporation.
- [21] Statement, F.P. (2009) Smart Grid Policy. Technical Report, Federal Energy Regulatory Commission.