

SUSTAINABILITY IN INFORMATION STEWARDSHIP:

TIME PREFERENCES, EXTERNALITIES, AND SOCIAL CO-ORDINATION

CHRISTOS IOANNIDIS, DAVID PYM, AND JULIAN WILLIAMS

ABSTRACT. The concept of stewardship in environmental economics is an established tool for environmental and natural resource management and the mitigation of risk from climate change. Similar concepts are well-established in accounting and management. Despite the ubiquity of the concept of stewardship, there is no generally accepted definition. We define the information steward as the agent/institution who enhances the system's resilience and sustainability, by maintaining and extending the life of its nominal operational capacity. Unlike individual agents who are not able to value systemic losses, the steward, whose function is the viability of the system as a whole, values such damages that degrade the system at a higher rate by adopting a lower discount rate. In the presence of deliberate attacks that degrade the information used/kept in the system, individual agents' defensive expenditure is always lower than the expenditure undertaken under instructions from the information steward. The ability of the steward to mobilize the totality of agents reduces the number of attackers, increasing the systems sustainability under a variety of technological considerations. The resulting configuration of defensive expenditure, although higher than the level that individual agents would have chosen based on their own valuation of their expected losses, ensures that the overall probability of successful attacks falls, significantly extending the system's sustainability.

1. INFORMATION STEWARDSHIP

The concept of stewardship in environmental economics is an established tool for environmental and natural resource management (e.g., [11]) and the mitigation of risk from climate change (e.g., [27]). Similar concepts are well-established in accounting (e.g., [14]), management (e.g., [9]), and insurance (e.g., [26, 5]). Despite the ubiquity of the concept of stewardship, there is no generally accepted definition.

Generically speaking, stewardship refers to the function of maintaining a given status of system. In the field of information security, the concept of stewardship has been used, albeit at a quite high level of abstraction, from the perspective of information management [24, 3]. The aim of this paper is to develop a more concrete definition of information stewardship and provide an economic model demonstrating the impact of the exercise of stewardship on the behaviour of information systems.

Inspired by the literature in environmental science, and bearing in mind existing work in information management, we define the role of the steward as the institution which maintains the system's resilience and sustainability, in the presence of unanticipated shocks that degrade nominal operating conditions.

By resilience (cf. [16]), we mean the system’s internal capacity to restore to an acceptable operating state. Consider the ‘marketplace’ ecosystem¹ of an internet retailer, whose IT system experiences a shock, such as major DDOS attack. The system will be deemed resilient if, after such an attack, to restore itself to its usual operating capacity by rapidly and progressively isolating the attackers and restoring services to legitimate users. In this case, the resilience of the system is characterized by the speed of service restoration and the quality of the restored service for legitimate customers. Organizations may differ in their preferences for rapidity and quality of service restoration.

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. However, 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 systems operational capacity to a steady, albeit lower, level. In both circumstances, the impact of the steward is to restore the system’s equilibrium path in the presence of disturbances, rendering the system resilient to them. It is the steward’s ability to reverse the divergent paths, harming all the agents in the ecosystem following shocks, which enhances the welfare of all its participants. All such mechanisms/procedures are put in place at t_0 ; that is, the steward anticipates the possibility of shocks and adopts the required policies in advance of their (shocks) realization. The steward therefore prepares the system to be resilient, rather than simply reacting to the shocks when they happen.

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. Consider again the internet retailer, whose role is to co-ordinate a ‘marketplace’ of providers of goods and services. By underwriting payments to both sides of this market, the market remains liquid and functioning. In the absence of such action and, given that in this market transactions are not supported by personal verification, the individuals’ incentive structure may lead market failure. In both such cases, the actions of the retailer maintains the system’s capacity by providing sufficient system resources and appropriate management policies and so acts as the steward of the ecosystem.

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 systems inability to perform within the acceptable bounds by t_1 . The steward’s contribution to

¹Following Moore [21], Nardi and O’Day [22] and others, we define an information ecosystem as an economic community of interacting (IT) networks, exchanging information and executing transactions according to agreed protocols, possibly under different jurisdictions. The underlying notion of system is described in terms of the concepts of distributed systems theory [8], and can be modelled mathematically as in [7].




FIGURE 1. Resilience

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 systems 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.

We can consider the stewardship problem as being analogous to the long-term asset management problem (sometimes called Merton's annuity problem). Over the very long term, risky assets grow, on average, at a faster rate than risk-free assets and therefore need to be included in the profile of assets. The balance of the portfolio, accounting for long-term risk preferences, will necessarily contain a subset of risky assets: this is necessary to ensure that the terminal value of the investment portfolio provides a suitable payoff (this is sustainability). However, when we address a shorter time horizon — that is, when considering the resilience problem — the investment planner is exposed to short term shocks that can be significant. The 'gamblers ruin' problem addresses whether the the investment planner should realize the loss and abandon an asset that underperforms in the short term or maintain the investment. For example, if a computer system is compromised by a security shock, the planner may choose to invest in order to restore the system's security or may choose to discontinue its use. If the portfolio has high levels of




FIGURE 2. Sustainability

diversification, then the need to abandon assets usually has a negligible impact on the wider portfolio. However, the cost of maintaining such a broad portfolio is very high and may result in the planner missing his terminal future wealth targets. The portfolio manager must therefore manage the combinations of risks within the portfolio, relative to the friction costs of diversification and any inter temporal requirements to liquidate under performing assets. The long-run planning problem is thus a sustainability argument: the reason to have the portfolio is that it can feasibly reach the terminal income targets required by investors. However, the short-horizon management requirements represent the need to manage the resilience of the portfolio (i.e., could the portfolio be subject to a foreseeable shock would reduce its ability to meet the terminal target toward zero and hence need to be fully liquidated). In this case, the policy problem is one dimensional and the parameters of the problem are fully exogenous, and are in effect captured by two variables: the slowly emergent mean drift of the portfolio and the more quickly realized volatility of the combination of assets, with controls dictated by the varying weights of the asset allocation.

There are prominent examples of the exercise of information stewardship in UK, EU, 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) encourages and maintains trade between the two economics. Both cases are examples of sustainability as the stewards 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 exercise of stewardship is costly to the agent as it requires investment of resources in things like infrastructure and monitoring. The motives for engaging are diverse. In the example of the retailer, above, its motive for engaging with the ecosystem as a steward can be legitimately perceived as profit maximization. However, this 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 [4, 2, 10]. 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 (e.g., privacy) to which their individual actions contribute, but also they will care for the consumption of this good by others (e.g., an agent may care its personal privacy, but also care about the privacy enjoyed by others). In the examples discussed above, the campaigns for the establishment of the UK's Data Protection and Freedom of Information Acts provide evidence of such behaviour. Second, agents wish to act as public benefactors because they derive satisfaction from doing so. For example, lawyers sometimes engage *pro bono* in protecting those whose privacy has been invaded.

The behaviour of agents interacting in a system of exchange is conditioned by their preferences, the pervading legal framework, and existing social conventions. Such conventions, known as norms, are either descriptive — that is, what agents in the system 'do' — or prescriptive, influencing what behaviour ought to be. The legal framework expresses society's values and determines the consequences (punishment) for actions deviating from such values. Benebou and Tirole [4] argue 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 incorrect beliefs regarding the contributions of the others and thus errors in their perception of the societal norm. The steward, by dispelling such misperceptions, can attain substantial benefits for the system as participants modify their behaviour. 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, the option available to him is 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 well-informed compared to individuals about the currently prevailing community standards. In 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 Funk [13], Tyran and Feld [28], and others. Such studies show that compliance is raised when its level has been chosen through a voting decision by the participants of the community/ecosystem. The law-maker/principal/steward, in setting the law and other obligations/incentives, must take into account the impact of his actions in the formulation of the norms which will be now expected to prevail. For example, we may consider that a steward signalling very high levels of IT defence expenditure 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. In our work, we adopt the convention that the steward is aware of the distribution of preferences across the agents of the ecosystem and that he is capable of enforcing the expenditure required on all agents who cannot leave the system. The norm in this case is the valuation of losses chosen by the steward. Agents must comply to remain in the system: their benefits are assumed constant and well above their expected losses, and they do not appear explicitly in the calculations.

In this paper, we consider the role of a steward whose only concern is the sustainability of the ecosystem in the presence of threats to information security. We assume that the steward is able to pass on the cost his decisions to the agents in the ecosystem whilst improving the agent's security. In this set-up, individual agents are assumed face a known probability of successful attack and they possess known mitigation technologies. To demonstrate the role of a steward in enhancing the sustainability of the ecosystem, we develop the following economic structure. In the first instance, we establish the equilibrium level of threat in the absence the steward (i.e., in a decentralized market, when agents make decisions based on their own constraints). We then compare the situation when a steward is present whose aim is to enhance the sustainability of the system. By allowing different pure time preferences, we account for the possible 'cultural' differences between the steward and the agents.

In Section 2, 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.² We introduce the concept of a steward who co-ordinates the defensive expenditure of the targets. We solve the model for three cases. In the first, we consider the general case in which the steward minimizes the total present value of the targets' expected losses using the individual targets' time preferences (along the lines suggested in [17]). We show that the presence of the steward is at least as effective as his absence. In the second, under mildly restrictive assumptions, the steward endogenizes attacker behaviour and sets the Pareto efficient allocation that minimizes the present value of the targets' expected losses. In the third, the steward, as in Figure 2, seeks to sustain the life of the ecosystem to at least the planning horizon by setting mandatory levels of defensive expenditure for each target. The steward sets these levels using time preferences for the valuation of expected future losses from successful attacks over the period of his planning horizon.

²It is straightforward to vary the model to account for payoffs that are utility-maximizing, but not necessarily monetary; for example, as in terrorism.

In Section 3, we present a discussion of how a steward’s time preferences emerge. We explore how, when the steward’s time preferences are not aligned with those of the targets (typically, when the steward’s time preferences are longer term than those of the targets), the allocation of mandatory defensive expenditure is no longer Pareto optimal. To motivate the choices of parameters taken in the three cases discussed above, we consider both evidence from the literature and an original dataset obtained from a survey Small-Medium Enterprises (SMEs) in the UK. Finally, we discuss briefly how steward’s might emerge in ecosystems of the kind we have considered.

We conclude the paper with a brief discussion of the policy implications of our work. We also discuss briefly some future directions for our work.

2. THE MODEL

Our starting assumption is 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. This stands in contrast to having a Nash policy-maker, who chooses a policy in equilibrium with the ecosystem participants. If the Stackelberg policy-maker is benevolent, he will tailor a policy so as to drive the participants towards the Pareto optimal allocation of resources.³ In our case, the allocation of resources is the chosen level of defensive expenditure by the participants under attack from an optimizing antagonist.

In our model, we make two simple structural assumptions. First, that increasing defensive expenditure by an ecosystem participant reduces the likelihood of successful attacks on that participant. Second, for any given participant, the greater the number of attackers, the higher the likelihood of a successful attack.

In [17], the present authors outline a framework for security investment decision-making that uses continuous discount rates (i.e., exponential discounting) to describe time-preferences. These preferences capture the trade-off between present expenditure and future losses from successful attacks. This representation renders the description of the mechanism of attack and defence as simple rate simple rate parameters. The usefulness of this algebraic representation is that the resulting loss functions are continuously differentiable.

Consider an ecosystem with N_T targets and N_A attackers. We set the ratio $\eta = \frac{N_A}{N_T}$; that is, the number of attackers per target. In general, we will consider the pool of attackers to be homogenous and N_T is always assumed to be fixed.

Fixing the number of attackers per target, η , we examine the choice of defensive expenditure x_i for the i^{th} target at time zero. 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. However, even with a single attacker attacking a large group of targets, given infinite time the probability of a successful attack, even from a single attacker amongst a large group of targets,

³For example, a Stackelberg policy-maker regulating a prisoner’s-dilemma game would impose dove-dove cooperation prior to the game’s commencement, anticipating the Nash equilibrium of the original game.

should converge to unity. Therefore, for any $\eta > 0$, $\lim_{t \rightarrow \infty} \Psi(\eta, \alpha_i, t) = 1$, a functional form that satisfies this condition 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 $t_0 < t \leq T$. Setting $\alpha_i = 0$, we see that $\Psi = 0$, so that the probability of a successful attack is zero for all time periods. In this situation, attackers have no ability to inflict losses on targets.

The probability law that drives the probability of successful attack favours an eventual successful attack over time. For instance, if $\alpha = 0.9$, then, if a firm is being attacked by a single attacker, the probability of a successful attack is 99.9% in 7.67 years in the absence of defensive expenditure, we shall assume that rate parameters are measured in per annum rates.

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 assumed to be fixed over $t_0 < t \leq T$. An alternative way of thinking of L_i is that if the present value of information assets is fixed to be M_i , then L_i is the value that satisfies $M_i = \int_{t_0}^T e^{-\beta_i t} L_i$, or

$$L_i = \frac{e^{T\beta_i} M_i \beta_i}{e^{T\beta_i} - 1}$$

for time horizon T . That is, we assume the current value of assets is evenly (by time discounted weighting) amortized over T . For certain types of firm, measuring the *instantaneous loss* L_i is quite easy. For instance, if a firm's assets are entirely information assets and it is financed by a mix of equity E_i and debt D_i , we can use the [20] model to determine the present value of assets M_i by assuming that the value of equity is a call option on assets with a strike price D_i . For reasonably long amortization, periods this formulation collapses to M_i tends to $D_i + E_i$.

In this case, L_i is maybe interpreted as the implied continuous dividend from information security assets amortizing over T . It is important to note that L_i is fixed for all $t < T$ and is measured in time t currency. Therefore, its present value is the discounted aggregation of L_i over $t_0 < t \leq T$.

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

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

The important point to note is that the attack likelihoods as a product of $(1 - e^{-\alpha_i \eta t}) e^{-\psi_i x}$ is an instantaneous probability. For instance, a single unit of defensive expenditure $x_i = 1$, when $\alpha = \psi = 0.50$ and a single attacker attacking the firm $\eta = 1$, the nominal (i.e. without discounting) factor multiplying the nominal loss L_i is 1.9192. When security investment is equal to ten units $x_i = 10$, then this factor reduces to 0.0213.

It is often helpful to measure the total nominal risk factor to one dollar of assets, when comparing time preferences this is simply imposing $L = 1$ and $\beta = 0$; that is,

$$R(x, \eta) = \int_{t_0}^T (1 - e^{-\alpha_i \eta t}) e^{-\psi_i x} dt,$$

where $R(x, \eta)$ is the total risk factor for a single unit of information assets at risk. By construction, this need not be less than unity as the same asset maybe at risk over the continuous interval t_0, T . $R(x, \eta)$ tends to a constant when $T \rightarrow \infty$ provided α and ψ are greater than zero.

2.1. Target Security Investment. We now assume that 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 [15].

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. For instance, if $\psi = 0.5$, then a single unit of defensive expenditure will reduce the probability of a successful attack by a factor of 0.6065 throughout the time interval t_0, T , ten units of present value defensive expenditure reduces the likelihood of attack by a factor of 0.0067.

The expected present value of losses is therefore

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

If 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 :

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

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

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

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.

2.2. The Market for Attacks. We now assume that 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

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

Targets are assumed to have knowledge of the average defensive expenditure $\tilde{x} = N_T^{-1} \sum_{i=1}^{N_T} x_i$, the average security decay factor $\tilde{\alpha} = N_T^{-1} \sum_{i=1}^{N_T} \alpha_i$ and the average security effectiveness factor $\tilde{\psi} = N_T^{-1} \sum_{i=1}^{N_T} \psi_i$. For ease of exposition, each expectation is supposed to be independent, and we assume that the signal extraction problem for attackers is relatively acute and therefore the anticipated rewards are also set in expectation $\tilde{R} = N_T^{-1} \sum_{i=1}^{N_T} R_i$. If attackers are randomly assigned to a particular target, then the time t reward from a single attack is

$$V(t) = N_A^{-1} N_T \tilde{R} (1 - e^{-\tilde{\alpha}t}) e^{-\tilde{\psi}\tilde{x}}.$$

Rewriting this equation in terms of η , we obtain

$$V(t) = \eta^{-1} \tilde{R} (1 - e^{-\tilde{\alpha}t}) e^{-\tilde{\psi}\tilde{x}}$$

and the present value of attacks is therefore

$$PV_A(t) = \int_{t_0}^T e^{-\gamma t} \eta^{-1} \tilde{R} (1 - e^{-\tilde{\alpha}t}) e^{-\tilde{\psi}\tilde{x}} dt.$$

Evaluating this integral, setting $t_0 = 0$, we obtain

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

The marginal attacker enters the market until the present value of expected rewards, $PV_A(T)$, equals the value of costs C_A . This is different from a single attacker making the decision to

engage in N_A multiple attacks as the decision function in this case would be in the form of a profit maximization problem rather than a binary market entry decision. In the 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. If attackers are risk neutral, the boundary condition for the marginal attacker choosing to enter the market for attacks equates the present value of an attack PV_A to cost of launching each attack:

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

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

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

For the first-winner-takes-all attackers, competition is acutely intense. However, attackers are assumed to value one present-value expected-dollar of gain the same as a certain-dollar of cost, substantially increasing the viability of market entry. In addition, over time the assumption of positive α_i s results in a tendency for the likelihood of a successful attack to converge to unity, again giving attackers a greater nominal chance of success per dollar of expenditure.

To state our first proposition, we need some remarks. First, for tractability of exposition — without much loss of generality — we simplify our set up to the case in which we have N_T ex-ante identical targets. That is, each of the parameters is constant over its indexing set: for all i , $R_i = R$, $\alpha_i = \alpha$, $\beta_i = \beta$, $\psi_i = \psi$, and $L_i = L$. Second, we assume the identical targets are so *ex ante*; that is, they are independently identical and not a representative 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:*

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

Proof. See the appendix. □

2.3. Introducing the Steward. We now introduce a steward whose sole objective is to improve the ecosystem's sustainability. The steward, assumed to be fully informed and so aware of all parameters, can impose his choice of defensive expenditure, x_i^P , on each individual target, i . Therefore, in reaching his choice, he takes into account its impact on the actions of attackers.

If, in valuing future losses, the steward adopts the time-value preferences for losses of the agents in the ecosystem — as represented by their discount rates, β_i — as his discount rates δ_i , then he acts as a conventional social planner, seeking to establish the Pareto efficient allocation of defensive expenditure.

The policy-maker also allows us to eliminate T as a variable from our model. If we assume that the smallest discount rate the policy-maker uses for their time preference is $\tilde{\delta} = \max\{\delta_i\}$, for $i \in 1, \dots, N_T$, then from elementary calculus we know that $1 = \int_0^\infty \tilde{\delta} e^{-\tilde{\delta}t} dt$.

Therefore, if we set $\lambda = \int_0^T \tilde{\delta}^{-1} e^{-\tilde{\delta}t} dt$, for a value λ arbitrarily close to 1, then we can rearrange and compute $T = \delta^{-1} \log(1 - \lambda)$. For instance, if $\lambda = 0.99$ or 99% of the present value and $\delta = 0.10$ or 10% per annum, then $T = 46.0517$ years; that is, the contribution of one dollar of losses in 69.0776 years is equivalent to 1 cent in the present.

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

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

where

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

The steward therefore anticipates his impact on the market for attacks as he sets x_i^P , rather than in the case of individual firms, where x_i^P is exogenous and set in equilibrium.

We now discuss three cases: first, the case in which $\delta_i = \beta_i$, for all $1 \leq i \leq N_T$; second, the case in which all targets are identical, $\beta_i = \beta$ and $\delta = \beta$; and, finally, the case in which all targets are identical and $\delta < \beta$.

The first case allows for a general discussion of the Nash equilibrium versus the Pareto optimal allocation given the choice of time preferences by the individual targets.

The second case 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.

In the final case, we relax the assumption that the steward identifies the discount rate of individual firms and sets a specific discount rate $\delta < \beta$, where the steward values future security outcomes more highly than an individual firm. In this case, we can identify a time horizon $T^P = \delta^{-1} \log(1 - \lambda)$, over which the steward places significant value to losses relative to $T^\beta = \beta^{-1} \log(1 - \lambda)$ for the individual firm.

Case 1: Heterogeneous targets and homogeneous attackers. We briefly review the case in which $\delta_i = \beta_i$. Analytic solutions are limited as we cannot derive specific solutions for x_i^P without distributional assumptions for L_i , β_i , ψ_i , and α_i . In this case, and in the next case, the steward acts as a benevolent social planner minimizing total present loss for all firms at their discount rates.

The first point is that x^P is derived by including for the impact of the policy rate of defensive expenditure for each target simultaneously. In contrast, the equilibrium level of defensive in the absence of the policy x_i^N is set in respect to the randomly assigned total attacking effort. Attackers will enter and exit the market in equilibrium until

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

recalling that we assume that γ is the same for all attackers. Setting the individual level of equilibrium defensive expenditure,

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

the aggregate defensive expenditure is $\sum_{i=1}^{N_T} x_i^N$ and the present value aggregate expected total loss TL across all targets is

$$TL^N = \sum_{i=1}^{N_t} \frac{L_i \left(\beta_i + e^{\alpha_i \eta T} (\alpha_i \eta (e^{\beta_i T} - 1) - \beta_i) \right) e^{-T(\alpha_i \eta + \beta_i) - x_i^N \psi_i}}{\beta_i (\alpha_i \eta + \beta_i)} + x_i^N$$

over the time horizon T . Where η^N is the equilibrium from solving η^* by substituting $\tilde{x}^N = \sum_{i=1}^{N_T} x_i^N$ for \tilde{x} . The steward's allocation of x_i^P for each target aggregates as

$$TL^P = \sum_{i=1}^{N_t} \frac{L_i \left(\beta_i + e^{\alpha_i \eta T} (\alpha_i \eta^P (e^{\beta_i T} - 1) - \beta_i) \right) e^{-T(\alpha_i \eta^P + \beta_i) - x_i^P \psi_i}}{\beta_i (\alpha_i \eta^P + \beta_i)} + x_i^P,$$

for $\tilde{x}^P = \sum_{i=1}^{N_T} x_i^P$. Rearranging, we obtain

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

The steward improves the ecosystem by finding a set of x_i^P , where the total loss is $TL^P < TL^N$. By inspection, we can see that the steward can, do at least as well as the Nash equilibrium in assigning allocations of defensive expenditures; that is, the case when $TL^P = TL^N$ occurs when $x_i^P = x_i^N$. Moreover, in the case where $\gamma > 0$ — that is, attackers have a time preference for near-term expected rewards versus far term rewards — we can see that η^P is declining in \tilde{x} . Therefore setting $x_i^P = x_i^N + \xi_i$ for a small increment ξ_i per target will result in $TL^P < TL^N$, when $\alpha_i, \beta_i, \psi_i$ and L_i are fixed. We can also average TL^P over each target N_T to get the loss per target $L^P = TL^P / N_T$. In the case of heterogeneous targets, this is simply a scaling; for homogenous targets, we will see that this is a useful metric.

Case 2: Homogeneous targets and homogeneous attackers. Exploration of heterogeneous targets case is hampered by the fact that the optimization problem to compute TL^P is N_T dimensional. This is in fact the only major argument against mandatory defensive expenditure in the presence of externalities. The identification of the optimal vector x_i^P maybe too difficult for a single planner. However, for exposition purposes simplifying to N_T ex-ante identical targets 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

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

The policy-maker then nests the attacker intensity directly into their optimization problem (instead of the solution being a simultaneous equation problem in the Nash equilibrium context).

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

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

This is, of course, the discounted present value using the policy-maker's discount rate in the first instance $\delta = \beta$. Differentiating with respect to x yields our second proposition.

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:*

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

Proof. A simple argument from statements above, setting $\frac{\partial L^P}{\partial x} = 0$. □

2.3.1. Discussion of Proposition 2. This equation is not analytically solvable for x^P in generality, but is a relatively easy to solve once the policy maker and attacker discount rates δ, γ are chosen (T is therefore defined in terms of $\max(\delta, \gamma)$ and the nuisance parameter $\lambda \rightarrow 1$) and the nominal loss L . The remaining terms are the technology parameters α and ψ which are subject to uncertainty.

We shall now explore an example where the social discount rate δ is equivalent to the private discount rate β .

2.3.2. Example UK SMEs versus Black-Hole Vulnerability Kit Hackers. For our first worked example we use a small survey (see Table 1, below) of UK SMEs to gauge the firms' discount rates and then explore the implications for changing the technology parameters α and ψ .

From a small survey of 59 SMEs, participating in the UK National Information Security Conference in 2010, 81% of respondents reported that information assets were amortized over 5–7 years. The implied discount rate for this is around 32%, if we assume 90% of the present value is contained in that timeframe.

All of the respondents reported that they used versions of Microsoft's Windows operating system. In fact, the majority reported that they used Windows XP, suggesting that a time horizon of greater than nine years for information assets is possible, so indicating that a private discount rate as low as $\beta = 0.25\%$ may also be realistic.

The firms indicated an average value of their information assets at risk to be around £2 Million, versus an average turnover of £5 Million. Therefore, we make the assumption that L (\approx £500,000). We state all losses in terms of expected loss per present value 1000s of pounds.

To obtain a point-estimate of γ , the attacker discount rate, we review the pricing of exploit kits and appeal to financial economics to deduce the correct discount rate of future expected rewards. Exploit kits are used to deliver malware to computer systems and are designed to exploit vulnerabilities in common operating systems such as Microsoft Windows. The exploit kit is a delivery system that is sold to attackers who design payloads to be deployed by the exploit kit. Variations in pricing for these kits can be used to determine the time value of this technology and hence the indicative future rewards from their use.

The ‘Black-Hole’ exploit kit as of 29 March, 2012 accounted for 28% of the delivery market for detected malware. The kit targets vulnerabilities in older versions of many common web browsers, including Microsoft’s Internet Explorer, Google’s Chrome and Apple’s Safari. Price variation in Black-Hole (see [1] indicates that year on year attackers are discounting future rewards at quite a low rate, circa 15%.

We therefore use $\delta = \beta = 25\%$ versus $\gamma = 15\%$ for our simulation. We set the cost-reward ratio for attackers to be $C_A/R = \log(1 - 0.15) \approx 0.85$, to match the attacker discount rate. If the market for attacks mimics a normal market operation then the discount rate on expected returns should match the opportunity cost, therefore $c = 1 - \gamma$, if attackers are truly risk neutral. Investigation of attacker costs and discounting is left for future empirical research.

We now set up the following simulations: set $\psi = \{0.1, 0.25, 0.5, \}$ (defence is not effective to very effective) and α is in the range $\{0, 2\}$ (attacking effort is not effective to very effective): we compute x^N versus X^P over the range of α . Next, we 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 3.

Case 3: Policy-maker with long-term time preferences. We now consider the example from Case 2, but we set $\delta < \beta$; 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 in Section 1.

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. For our empirical analysis, we assume that δ is now 10% rather than 25%. In this case, the policy-maker time horizon extends from 9.2103 to 23.0259 years if we consider $\lambda = 90\%$ of the present value.⁴

⁴The half-life of the present value has now extended from 2.7726 to 6.9315 if we consider 50% of the present value.

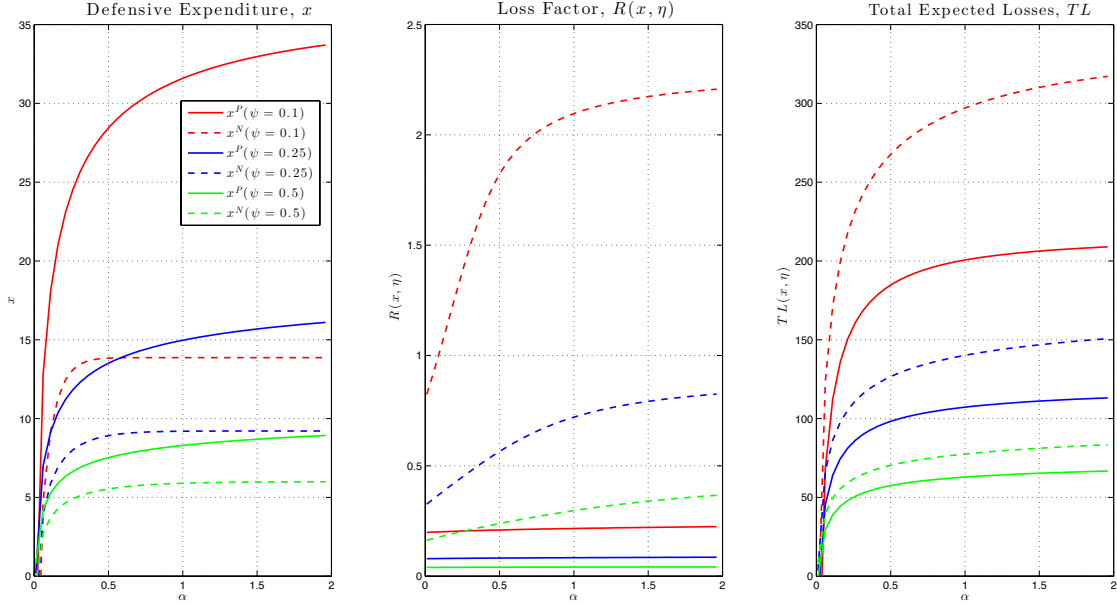


FIGURE 3. Comparison of the impact of a steward on defensive expenditure, risk, and expected losses assuming that the ecosystem consists of UK SMEs of the type surveyed in Case 2. For tractability, we assume identical targets. The level of defensive expenditure for each target is denoted by x (the left plot), the expected loss factor $R(x, \eta) = \int_0^T (1 - e^{-\alpha\eta t}) e^{-\psi x} dt$ (centre plot) and total expected losses $TL = \int_0^T e^{-\delta t} L (1 - e^{-\alpha\eta t}) e^{-\psi x} dt + x$ (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 $\int_0^T e^{-\beta t} L dt = \text{£}2 \text{ Million to } \text{£}555,555 \approx \text{£}500,000$ and TL is presented per $\text{£}1000$ of assets. Attacker rewards are set such that they receive 0.15 units of revenue per unit of effort therefore $c = 0.85$

Therefore, for any t in $t_0 < t < T$, the policy-maker values expected losses more highly than the ex-ante identical targets and sets levels of x^P accordingly. In this instance, the level x^P from the view of the targets is not the Pareto optimal allocation, which occurs when $\delta = \beta$. We denote the level of defensive expenditure allocated by a policy-maker when $\delta = \beta$ by x^P and denote the case when $\delta < \beta$ by x^S .

Using the previous case as a starting point, we can see that for a steward solving for x^P versus x^S , the minimum loss shifts to the right; see Figure 4, for the steward's switching from $\delta = 0.25$ to $\delta = 0.10$. The steward now weights more of the loss from expected attacks versus the immediate expenditure on defensive action. An interesting point to note is that the steward cannot achieve the lowest present value of losses. For instance, in this case for $\delta = 0.25$, the

TABLE 1. The following survey (59 respondents to questionnaire) was carried out using the SNAP survey software tool. Results are abridged for UK SMEs from the NISC May 2010 meeting. Results are for indicative use only, the sample is a self-selecting group of Chief Information Security Officers. The survey followed up a series of structured interviews with participants at this meeting. Full survey instrument and results are available from the authors. The columns have the following definitions: Turnover (expected annual turnover in coming year); Amortization (length of time amortize information assets); Standard (do you hold a standards certification?); Assets (estimated value of current information assets). All values are in £.

Turnover	Results	Amortization	Results	Standard	Results	Assets	Results
< 500K	6 (10.2%)	> 0 - < 2 yrs	2 (3.4%)	National	24 (40.7%)	< 250K	11 (18.6%)
> 500K-<2M	23 (39.0%)	> 2 - < 5 yrs	5 (8.5%)	International	31 (52.5%)	> 250K - < 1M	18 (30.6%)
>2M-<5M	30 (50.8%)	>5 - < 7 yrs	48 (81.4%)	None	4 (6.8%)	> 1M - < 2M	26 (44.0%)
> 5M	0 (0.0%)	> 7 yrs	4 (6.7%)			> 2M	4 (6.8%)

lowest loss is under 9 units of present value pounds per thousand present value at risk, whereas when $\delta = 0.10$ the lowest attainable loss is over 9.5 pounds per thousand at risk.

This does not indicate that the policy-maker is performing in a worse fashion when $\delta = 0.10$, simply that he values future losses substantially longer into the future when the likelihood of successful attack (mediated by $\alpha \times \eta$) is substantially higher. By construction, if η was fixed exogenously and $\alpha \times \eta = \delta$, then the present value of total losses would be identical for all δ at the optimal loss point.

We now 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 with the same properties in Case 2; that is, $L = 500,000$, $\gamma = 0.15$, $c = 0.85$ and $\beta = 0.25$. In this case, we assume that the steward has adjusted his time horizon from 9 years to 23 years covering 90% of the present value (or from 2.7 years to 7 years for the half life of the ecosystem). Figure 5 presents the variation in defensive expenditure x , loss factor $R(x, \eta)$, and total loss $TL(x, \eta)$ for the Pareto optimal case x^P versus the longer term steward x^S .

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. However, the downside from the steward's longer term time preferences is illustrated in the third panel of Figure 5. 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 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). We see that the externality created by the attacker choice η interacts with the new choice of δ . The steward views the externality as being larger than the individual targets (as the attacker




FIGURE 4. Total expected loss versus mandatory investment x , where the minimum point is denoted x^P , when $\delta = \beta$ (unbroken line) and $\delta < \beta$ (broken line), where $\delta = 0.10$. We label this new minimum x^S . Here we assume that $\psi = 0.25$, $\alpha = 0.25$, $\gamma = 0.15$, $c = 0.85$, and $\int_0^T e^{-\beta t} L dt = \text{£}2 \text{ Million}$.

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.

In the next section, we discuss the current evidence for the variation in time preferences for firms versus policy-makers.

3. EVIDENCE FOR THE TIME PREFERENCES OF THE STEWARD

In public economics, the steward's chosen discount rate δ is often referred to as the *social discount rate* or policy-maker time preference. The most commonly encountered public policy social discount rate is the base or policy rate dictated by central banks issuing fiat money. However, for most policy decisions requiring time preferences to be imposed either in legal structures or public investment- the central bank policy rate is not used as it is deemed to be a basic rate, and therefore a premium is added or subtracted.

The fact that private discount rates diverge from social discount rates is the subject of extended discussion in the economics literature (e.g., [6, 19]). However, the broad consensus for firm private discount rates focuses on the financial economic viewpoint. Models such as the capital asset pricing model indicate that log linear preferences relative to a single representative risky rate

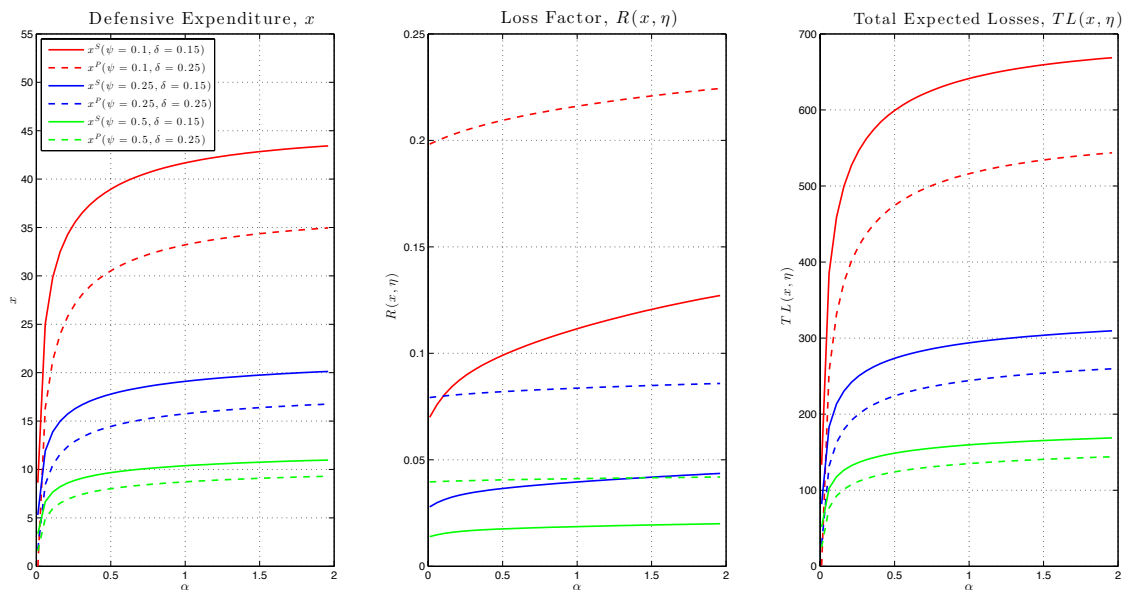


FIGURE 5. We compare the allocation of optimal investment x^P versus x^S , for the cases in which $\delta = \beta$ and $\delta < \beta$, respectively. The left hand plot presents the optimal allocation of defensive expenditure when $\delta = 0.25$ (dashed line) versus $\delta = 0.15$ (unbroken line) 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 $\int_0^T e^{-\beta t} L dt = \text{£}2 \text{ Million} \approx 500,000$ and TL is presented per $\text{£}1000$ of assets. Attacker rewards are set such that they receive 0.15 units of revenue per unit of effort; therefore, $c = 0.85$.

can reduce the discount rate problem for private firms to a simple measurement of the covariance of asset valuations to the broader economic system. The fact that a risk premium exists for firms indicates that if targets are assumed to be firms and the steward is a public policy-maker, then δ will be required to be less than β as risk premiums are always positive.

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. In the three cases in Section 2, we derived the firm-specific discount rates to be the rates that amortized their 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 his 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 costs are imposed at rates that individual participants in the ecosystem may believe to be unjust given their own time preferences.

The very low social discount rate problem is an area of active debate in environmental economics and in particular the economics of climate change. The UK Government's Stern Review [27] sets time preferences with respect to a discount rate approaching zero. This has sparked substantial debate in the economics literature, as future losses from climate change impacts have not been discounted at rates markedly similar to public or private investments; see [23, 29] for extended discussion. The issue is more acute here as losses from climate risks are generally assumed to be realized at a reasonable distance into the future. Therefore, even small discount rates have very little impact on the current cost-benefit analysis assessing risk mitigation.

For information security contexts, the impact of the time preference assumption is not so acute as investment horizons are much shorter (see for instance [17] for a model of investment horizons). However, the interaction of the externality with the differentiated discount rate between targets and the steward does indicate that this is an important issue for information ecosystems. Under certain measurements, targets may have very large discount rates amortizing information assets over periods as short as 12 to 18 months. In this case, even for the most benevolent technology risk environment ($\alpha \rightarrow 0$, $\psi \gg 0$), the attacker externality will need to be redistributed via mandatory policy interventions.

Our model presents the most basic externality case and does not include externalities created by costs directly attributed to the mechanism of attack. If, however, considerations of sustainability are not met by the social planning choice (i.e., $\delta_i = \beta_i$), then the steward adopts a more severe valuation of future losses in excess of the valuation assigned by individual agents. In effect, he imposes a uniform social discount rate δ over all targets i . Consider the case of identical targets (i.e., a uniform $\beta = \beta_i$, for all i). If $\delta = \beta$, 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 [6], 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. However, the speed of discounting by firms of their information security assets can be very high (see Table 1) suggesting that the rate could be as high as 40% per annum. It is unclear whether this discount rate also applies to future losses. If so, then the private discount rates would be expected to be very different from stated public discount rates that are normatively closer to 10% for developed countries (see US office of Management and Budget policy overview [19]).

Data such as that which we obtained for a sample of UK SMEs can be used to estimate, in the terms of Figure 2, the time t_1 (which is a function of x^N). Thus the steward can estimate his need to intervene in order to sustain the ecosystem within which the SMEs operate to the horizon T . The cost of such a potential intervention will be $x^S - x^N$, so he can determine the marginal increase in the time horizon with respect to additional investment.

4. CONCLUDING REMARKS

We have defined the information steward as the agent/institution who enhances the system's resilience and sustainability, by maintaining and extending the life of its nominal operational capacity. We have given a model that explores the case of sustainability in detail. The novelty of the model lies in the interaction, in the presence of externalities, between the time preferences of the steward and those of the participants in the ecosystem. We show that when the steward's time preferences are substantially longer term than those of the participants, the effect of the externalities is greatly magnified.

Future work will be to study models of resilience that are compatible with model of sustainability presented here (recall the discussion of the analogy with financial economics in Section 1). Modelling resilience will explain how stewardship institutions emerge as natural phenomena in decentralized systems in which all agents are experiencing recurring interactions and, as a result, patterns of behaviour materialize which establish desirable conduct, which is encouraged and policed.

In keeping with similarly specified models — such as [18] for pricing insurance contracts with endogenous attackers and [12, 15, 25] for defence expenditure allocation — our model is specified by two parameters, one each for the marginal effectiveness of the attackers and defenders. The need for stewardship, be it public policy-making (e.g., government) or ecosystem management (e.g., in Amazon Marketplace) intervention, is determined by these two technological parameters, both of which are difficult to calibrate. By introducing time preferences for future losses that are differentiated between the steward and the targets, we move to a situation in which we are able to explore the viable range of these parameters. In this situation, the steward allocates costs derived from externalities fairly amongst the targets without the need to identify these technological parameters with precision.

In the absence of a clear theoretical analysis to-date, little progress has been made in implementing effective regulatory régimes. Our model unambiguously demonstrates that for all plausible technology configurations the presence of a steward, as defined in this paper, is beneficial to the sustainability of the ecosystem.

Acknowledgements. We are grateful to Joe Swierzbinski for his general advice on this paper. We are also grateful to Adrian Baldwin, Iffat Gheyas, Bruce Hallas, John Morrison, and Simon Shiu for their assistance with empirical studies referred to in Sections 2 and 3.

REFERENCES

- [1] L. Allodi. The dark side of vulnerability exploitation. In G. Barthe and B. Livshits, editors, *Proc. International Symposium on Engineering Secure Software and Systems*, 2012. http://ceur-ws.org/Vol-834/paper12_essosds2012.pdf.
- [2] J. Andreoni. Giving with Impure Altruism: Applications to Charity and Ricardian Equivalence. *Journal of Political Economy*, 97(6):1447–1458, 1989.
- [3] A. Baldwin, D. Pym, M. Sadler, and S. Shiu. Information stewardship in cloud ecosystems: Towards models, economics, and delivery. In *Proc. 2011 Third IEEE International Conference on Cloud Computing Technology and Science (CloudCom 2011)*, pages 784–791. IEEE Digital Library, 2011.
- [4] R. Benabou and J. Tirole. Laws and norms, 2012. Working Paper IZA DP no 6290.

- [5] R. Böhme and G. Schwartz. Modeling cyber-insurance: Towards a unifying framework. In T. Moore, D. Pym, and C. Ioannidis, editors, *Economics of Information Security and Privacy*. Springer, 2010. Proceedings WEIS 2009, London.
- [6] A. Caplin and J. Leahy. The social discount rate. *Journal of Political Economy*, 112(6):1257–1268, 2004.
- [7] M. Collinson, B. Monahan, and D. Pym. *A Discipline of Mathematical Systems Modelling*. College Publications, 2012.
- [8] G. Coulouris, J. Dollimore, and T. Kindberg. *Distributed Systems: Concepts and Design*. Addison Wesley; 3rd edition, 2000.
- [9] J. David, F. Schoorman, and L. Donaldson. Toward a Stewardship of Management. *Academy of Management Review*, 22(1), 1997.
- [10] E. Deci. *Intrinsic Motivation in Human Behavior*. Plenum, 1985.
- [11] F.S. Chapin III, G. P. Kofinas, and C. Folke, editors. *Principles of Ecosystem Stewardship: Resilience-Based Natural Resource Management in a Changing World*. Springer-Verlag, 2009.
- [12] N. Fultz and J. Grossklags. Blue versus red: Towards a model of distributed security attacks. In R. Dingle-dine and P. 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.
- [13] P. Funk. Is there an Expressive Function of Law? An Empirical Analysis of Voting Laws with Symbolic Fines. *American Economic Review*, 9(1):135–139, 2007.
- [14] F. Gjesdal. Accounting for Stewardship. *Journal of Accounting Research*, 19(1):208–231, 1981.
- [15] L. Gordon and M. Loeb. The Economics of Information Security Investment. *ACM Transactions on Information and Systems Security*, 5(4):438–457, 2002.
- [16] C. Hall, R. Anderson, R. Clayton, , E. Ouzounis, and P. Trimintzios. Resilience of the Internet Interconnection Ecosystem. In B. Schneier, editor, *Economics of Information Security and Privacy III*, pages 119–148. Springer, 2013.
- [17] C. Ioannidis, D. Pym, and J. Williams. Fixed costs, investment rigidities, and risk aversion in information security: A utility-theoretic approach. In B. Schneier, editor, *Economics of Security and Privacy III*. Springer, 2012. Proceedings of the 2011 Workshop on the Economics of Information Security.
- [18] H. Kunreuther and G. Heal. Interdependent Security. *Journal of Risk and Security*, 26(2–3):231–249, 2003.
- [19] J. J. Lew. Discount rates for cost-effectiveness, lease purchase and related analyses. Technical Report OMB Circular No. A-94, 2011.
- [20] R. Merton. On the pricing of corporate debt: The risk structure of interest rates. *Journal of Finance*, 29(2):449–470, 1974.
- [21] J. Moore. Predators and Prey: A New Ecology of Competition. *Harvard Business Review*, 71(3):75–86, 1993.
- [22] B. Nardi and V. O’Day. *Information ecologies*. MIT Press, 1999.
- [23] W. D. Nordhaus. The “Stern Review” on the Economics of Climate Change. Technical Report w12741, NBER Working Paper, 2007.
- [24] D. Pym and M. Sadler. Information Stewardship in Cloud Computing. *International Journal of Service Service, Management, Engineering, and Technology*, 1(1):50–67, 2010.
- [25] D. Pym, J. Swierzbinski, and J. Williams. The need for public policy interventions in information security. Manuscript at <http://homepages.abdn.ac.uk/d.j.pym/pages/InfoSecPubPol.pdf>. Submitted for publication.
- [26] N. Shetty, G. Schwartz, M. Felgyhazi, and J. Walrand. Competitive cyber-insurance and internet security. In T. Moore, D. Pym, and C. Ioannidis, editors, *Economics of Information Security and Privacy*, pages 229–248. Springer, 2010. Proceedings WEIS 2009, London.
- [27] N. Stern. *Stern Review on the Economics of Climate Change: Executive Summary Long*. HM Treasury Stationary Office, 2006.
- [28] J. Tyran and L. Feld. Achieving Compliance When Legal Sanctions are Non-Deterrent. *Scandinavian Journal of Economics*, 108(1):135–156, 2006.
- [29] M. L. Weitzman. A Review of the Stern Review on the Economics of Climate Change. *Journal of Economic Literature*, 45(3):703–724, September 2007.

APPENDIX A. PROOF OF PROPOSITION 1

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:

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

Proof. We know that for the targets $1 \leq i \leq N_T$, the optimal level of defensive expenditure is solved by

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

where β_i is the target's discount rate, α_i is the security decay factor, ψ_i is the security effectiveness factor, η is the number of attackers per target, and T is the time horizon.

Setting each firm to be ex-ante identical — that is, $R_i = R$, $\alpha_i = \alpha$, $\beta_i = \beta$, $\psi_i = \psi$, and $L_i = L$ — the choice of defensive expenditure x_i for each firm is therefore dependent on the attacking intensity directed against that firm.

If attackers randomly choose firms to attack — that is, the index of the chosen firm to be attacked is uniformly distributed — then the total number of attackers N_A is distributed evenly amongst the N_T targets. Therefore, $\eta = N_A/N_T$ is identical for all targets and so all targets will choose the same defensive expenditure $x = x_i$.

The marginal attacker (who observes the existing set of attackers per target and then decides whether to spend C_A to attack; he can choose to attack just once) observes that, for all targets, $\tilde{x} = N_T^{-1} \sum_{i=1}^{N_T} x_i$. Given that $x_i = x$, this simplifies to $\tilde{x} = N_T^{-1} N_T x = x$. In this case, the attacks identify precisely the level of expenditure chosen by the target to which they are randomly assigned. The marginal attacker also observes precisely $\alpha = \tilde{\alpha}$ and $\psi = \tilde{\psi}$.

Therefore the marginal attacker solves his entry decision-making problem

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

Dividing both sides of this inequality by R , and setting $C_A/R = c$, we find

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

The identical marginal attackers will enter the market until

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

Similarly, when

$$\frac{(\gamma + \alpha e^{T(\alpha+\gamma)} - (\alpha + \gamma)e^{\alpha T}) e^{-T(\alpha+\gamma)-x\psi}}{\gamma\eta(\tilde{\alpha} + \gamma)} < c,$$

attackers will exit the market until the number of attackers per target satisfies the following:

$$(1) \quad \eta^* = \frac{(\gamma + \alpha e^{T(\alpha+\gamma)} - \gamma e^{\alpha T} - \alpha e^{\alpha T}) e^{-T(\alpha+\gamma)-x\psi}}{c\gamma(\alpha + \gamma)}.$$

We assume that the number of attackers and the level of defensive expenditure now adjusts dynamically as attackers and targets approach equilibrium along their own reaction functions. Targets are ex-ante identical, and so each target solves

$$(2) \quad x^* = \psi^{-1} \log \left(\beta^{-1} L\psi e^{-\beta T} - (\alpha\eta + \beta)^{-1} L\psi e^{-T(\beta+\alpha\eta)} \right).$$

Replacing x with x^N and η^* with η^N in Equation 1 and replacing η with η^N and x^* with x^N in Equation 2, we obtain

$$\begin{aligned} \eta^N &= \frac{(\gamma + \alpha e^{T(\alpha+\gamma)} - \gamma e^{\alpha T} - \alpha e^{\alpha T}) e^{-T(\alpha+\gamma)-x^N\psi}}{c\gamma(\alpha + \gamma)} \\ x^N &= \psi^{-1} \log \left(\beta^{-1} L\psi e^{-\beta T} - (\alpha\eta^N + \beta)^{-1} L\psi e^{-T(\beta+\alpha\eta^N)} \right) \end{aligned}$$

as required. □

UNIVERSITY OF BATH

E-mail address: c.ioannidis@bath.ac.uk

UNIVERSITY OF ABERDEEN

E-mail address: d.j.pym@abdn.ac.uk

UNIVERSITY OF ABERDEEN

E-mail address: julian.williams@abdn.ac.uk