The Need For Public Policy Interventions in Information Security

David Pym, Joe Swierzbinski, and Julian Williams

Abstract Should public policy-makers set minimum levels of behaviour for individuals and corporations regarding information security policies and investments? We consider a model in which a finite number of targets are at risk of attack, attacks are costly, and have a finite probability of success. One important innovation is an explicit model of the decisions of potential attackers on whether to mount attacks. The model shows how the behaviour of attackers and the nature of the technological environment can create a role for a policy-maker to coordinate optimal minimum levels of protective expenditure for firms.

1 Introduction and Related Work

There is currently an ongoing policy debate concerning the appropriate nature and extent of regulation to maintain the security of information assets. On one side of the debate, some policy-makers argue that the provision of advisory information by

D. Pym
Computing Science, University of Aberdeen, King's College, Aberdeen AB24 3UE, UK. Part of this work was carried out whilst Pym was employed at HP Labs, Bristol, UK. e-mail: d.j.pym@abdn.ac.uk

J. Swierzbinksi (Corresponding Author)
Business School, University of Aberdeen, King's College, Aberdeen AB24 3QY, UK. e-mail: j.swierzbinski@abdn.ac.uk

J. Williams
Business School, University of Aberdeen, King's College, Aberdeen AB24 3QY, UK. e-mail: julian.williams@abdn.ac.uk

The authors would like to acknowledge the support of the Technology Strategy Board funded project "Cloud Stewardship Economics". We also thank Adrian Baldwin, Andrea Patacconi, and Simon Shiu for helpful discussions in the preparation of this paper.

governments and the use of voluntary standards would be the best choice. Others maintain that some form of compulsory regulation is required.

For example, the UK Parliamentary Office of Science and Technology made the following observation concerning the issue of cyber security. “Opinion is divided as to whether cyber security regulation by government would be the best way forwards. Regulation could increase the level of adherence to best practice, however it will always lag behind developments in technology and would be difficult to monitor.” (see POSTNOTE 389, Sept. 2011).

In the European Union, ENISA (the European Network and Information Security Agency) is seriously considering the use of compliance-based regulations to supplement voluntary approaches. (ENISA, 2012). In contrast, in the United States, the presumption is that what is required is a culture of voluntary good practice based on efficient information sharing. (See, for example, Blueprint for a Secure Cyber Future, Dept. of Homeland Security, 2011.)

Our paper demonstrates how a need for compulsory regulation can arise in the context of cyber security. We identify a taxonomy of externalities and analyze how these factors can cause a divergence between unregulated private actions and those that would minimize the overall cost to society from cyber threats.

In order to make our arguments more precise, we develop a formal model of attacker behaviour and the incentives of targets to invest in defensive expenditure. One important innovation of our model is to highlight the importance of the role of attacker behaviour in creating a need for policy intervention.

The model shows the circumstances under which the social and private incentives to invest in cyber security can be expected to differ. One surprising result is that even if the technological environment were to be modified to minimize conflicts between social and private incentives, the nature of attacker behaviour is itself likely to create an incentive for underinvestment in security and a consequent need for government regulation.

In the context of our paper, an externality exists when the defensive action of one prospective target of cyber attacks also affects the risks faced or losses incurred by other targets. Reductions in the risks or losses to other agents resulting from one target’s defensive expenditure are a valuable byproduct from a social perspective. However, it is reasonable to expect that an individual decision maker will tend to undervalue such beneficial effects for others when weighing the costs versus the benefits of additional defensive expenditure. This in turn will typically result in an underinvestment in security from the perspective of society as a whole.

Externalities can occur because of the nature of the technological environment. For example, a vulnerability in the software or hardware of one prospective target may create an avenue for attacks on other targets. In addition, when a successful attack is carried out on one target, this may create losses for others as well. For example, when one firm’s service is interrupted due to an attack, the firm’s customers may suffer losses that are not fully compensated by the firm in question. We will sometimes refer to externalities that result from the technological environment in which targets operate as “ecosystem” externalities.

The analysis in our paper demonstrates that the behaviour of attackers can be another source of externalities which create a role for public policy. The defensive expenditure of any one target can make attacks on the entire population of targets less attractive for potential attackers. By doing so, the defensive expenditure of one target can reduce the level of attacks on all targets, thus providing an external benefit.

There is evidence from documented attacks to suggest that attackers must balance expected rewards and costs in making attacks. For example, in 2011 an FBI indictment of six Estonian nationals on "click fraud" reported that approximately $14 Million was raised from the infection of four million machines using DNSChanger Malware.[1] However, this attack required substantial investment in legitimate web businesses as a front for the fraud.

There is also evidence that attackers dynamically readjust their effort in response to the behaviour of attackers and potential targets. For example, Herley (2012) observe that one reason for email phishing attacks is to identify the email users who are most likely to fall prey to an attack. They observe that such behaviour is only sensible if an attack on all potential targets is too costly and attackers intend to focus attacks on more vulnerable targets. Baldwin et al. (2012) show that spikes in attacks on specific systems can lead to mutual excitement of attacks on other systems. Such behaviour suggests that attackers respond to an indication of a profitable opportunity (i.e., the initial attacks) by launching more attacks.

The "conficker" computer worm provides an example of the importance of externalities in the context of cyber security. The conficker worm was first detected in 2008 and, at its peak in early 2009, had infected between 9 and 15 million computers.[2]

An interesting aspect of the conficker worm was that it posed relatively little danger to an individual infected machine, but turned this machine (usually in its down time) into a component of a larger "botnet" which was then used to mount attacks on larger computer systems via spam emails or denial of service style attacks.

One of the issues with combating such a worm was that many of the computers involved were commercial units housed in call centers and other large offices.[3] Because of the relatively small level of damage to individual machines and the relatively high cost of defending against the conficker virus, the time taken to mitigate this worm was relatively slow. Nearly four years after the worm's release, 1.7 million machines were still infected with the conficker worm.[4]

Although the conficker worm was not particularly sophisticated technologically, it exploited in a sophisticated way the perverse economic incentives created by externalities. The cost of mitigation for large offices was higher than the risk adjusted

[1] Source: FBI Press Release, Malware 110911 "International Cyber Ring That Infected Millions of Computers Dismantled", September 11, 2011.

[2] See: Markoff, John (2009-01-22). "Worm Infects Millions of Computers Worldwide". New York Times.

[3] See: Leffall, Jabulani (2009-01-15). "Conficker worm still wreaking havoc on Windows systems". Government Computer News.

[4] See: Microsoft Security Intelligence Report: Volume 11, Microsoft, 2011

cost to these offices of having the worm on their systems. Hence, many firms were slow to take action to remove it.

There is a large literature in Economics on externalities and related topics such as public goods. See, for example, Varian (2010) for an introductory discussion and Cornes and Sandler (1996) and Laffont (2008) for more advanced treatments. However, there have been relatively few applications of the economic theory of externalities to the field of computer security. Varian (2004) is one well known example that considers how the nature of technological externalities may affect the level of investment in the reliability of information systems.

An early contribution to the literature on investment in information security by Straub and Welke (1998) outlines a model of threat and countermeasure that models risk as a combination of attacker and defender effort. Treating risk as a function of defender effort, Gordon and Loeb (2002) present a model of decreasing marginal returns to security investment. They propose a residual risk function that relates investment to the probability of a successful attack. Optimal investment in security is, therefore, a tradeoff between the risk adjusted expected loss and the deterministic level of investment. Other threat models, such as Ioannidis et al. (2009, 2011, 2012); Chen et al. (2011) or Gordon et al. (2010), utilize a real options or portfolio optimization approach to model the defensive response of a firm. In the above papers, the behaviour of attackers is assumed to be exogenous in the sense that attackers do not respond to targets' actions.

Some papers consider the interactions between attackers and defenders. For example, in Cavusoglu et al. (2008) a firm's security manager must estimate an attacker effort function in order to compute the firm's optimal expenditure on security. Florencio and Herley (2011) consider the relationship between the incentives of attackers to mount attacks and the observed volume of attacks. The papers by Cremonini and Nizovtsev (2010) and Fultz and Grossklags (2009) model the level of security in a computer network as the outcome of a strategic game between attackers and defenders. To the best of our knowledge, our paper is the first to focus on the externalities which affect the strategic interaction between attackers and defenders and how these externalities create a role for public policy that is likely to involve compulsory regulation.

The paper now proceeds as follows, Section 2 outlines our model for the general case. Section 3 introduces a public policy-maker with a weighted average objective function and solves for the optimal policy. Section 4 provides an example of how to implement this framework and illustrates a complementary approach to demonstrating attacker effort as an externality. Finally, Section 5 offers some concluding remarks and ideas for further extension of this research area.

2 The Model

2.1 Technology of Attack and Defence

We consider a model in which a fixed number of targets $i = 1,\dots,N_T$ are at risk of attack. Let n_{Ai} denote the average number of attacks made against target i in any given period. Although we consider the case where n_{Ai} is specified exogenously as a benchmark, in general n_{Ai} is determined endogenously through our model of attacker behaviour.

Targets can reduce the probability that an attack is successful by engaging in defensive expenditure. Let $\mathbf{X} = (x_1,\dots,x_{N_T})$ denote the vector of defensive expenditures chosen by all the targets.

Let $\sigma_i = \sigma_i(\mathbf{X}, n_{Ai})$ denote the conditional probability that one or more attacks mounted against target i are successful. For the purposes of this paper, the technology of attack and defence is summarized by the properties of the functions $\sigma_i(\mathbf{X}, n_{Ai})$.

It is plausible that the probability σ_i should depend on the level of defensive expenditure by target i and the average number of attacks mounted against target i. Our model also allows the possibility that σ_i depends positively on the defensive expenditure of one or more other targets j. We say in this case that a technological or "ecosystem" externality exists, since the defensive expenditure of target j provides an "external" benefit to target i, and this benefit occurs because of the nature of the technological interactions that determine the environment or "ecosystem" in which attacks occur.

We assume the following properties of the functions $\sigma_i(\mathbf{X}, n_{Ai})$:

Property 1: $\partial\sigma_i/\partial n_{Ai} > 0$ for all i, so that an increase in the average number of attacks against a target increases the probability that at least one attack is successful for all levels of defensive expenditure x_i and any number of attacks n_{Ai};

Property 2: $\partial\sigma_i/\partial x_i < 0$ for all i so that an increase in the defensive expenditure of a target reduces the probability that some attack on that target is successful;

Property 3: $\partial^2\sigma_i/\partial x_i^2 > 0$ for all i, at least for large enough values of x_i. Property 3 implies that the marginal returns to defensive expenditure are decreasing, at least for large enough values of x_i;

Property 4: $\partial\sigma_i/\partial x_j \leq 0$, for all i and $j \neq i$. Property 4 ensures that, if an ecosystem externality exists, it represents a potential benefit in that firm j's expenditure reduces firm i's risk.[5]

[5] Note that ecosystem externalities need not be symmetric. It could be the case, for example, that $\partial\sigma_i/\partial x_j < 0$, but $\partial\sigma_j/\partial x_i = 0$. Moreover, ecosystem externalities may exist between some targets, all targets, or no targets.

2.2 Attacker Behaviour

For simplicity, we assume that all potential attackers are identical. In particular, each attacker has the same, constant cost, C_A, of mounting an attack. Suppose, in addition, that each attacker participating in attacks mounts one attack in any given period.

Let $R_i(n_{Ai})$ denote the expected monetary reward per attack against target i when one or more of these attacks turns out to be successful. We suppose that $dR_i/dn_{Ai} \leq 0$ to indicate the possible effects that competition among attackers to obtain a greater share of the reward can have on the expected reward per attack.

In order to highlight the effects of competition among attackers, we consider, in most of this paper, a version of the model where attackers obtain a constant expected reward R_i when there are one or more successful attacks on target i. Each attacker participating in attacks on target i is assumed to obtain an equal share of this reward. Hence, $R_i(n_{Ai}) = R_i/n_{Ai}$ in this case.

One circumstance that produces an equal share rule like that described in the previous paragraph is a case where the "first-winner-takes-all". In such a case, the first attacker who mounts a successful attack against target i receives the entire reward R_i, and all other attackers mounting attacks against this target receive nothing. If each attack has an equal chance of being the first one to be successful, then the probability that a given attack is the one to obtain the reward from "success" is simply $1/n_{Ai}$ and $R_i(n_{Ai}) = R_i/n_{Ai}$.

Other models of how attacks generate rewards are possible. At the other extreme from the "first-winner-takes-all" model, we might assume that successful attacks resemble an attack by n_{Ai} sharks against a shoal of fish. Each shark eats approximately the same chunk of fish from the school. Such a "shark attack" could be modeled by assuming that $R_i(n_{Ai}) = R_i$ for all n_{Ai}.

We suppose that attackers wish to maximize their expected profit. The expected profit which an attacker obtains from mounting an attack on target i is given by the following expression.

$$\sigma_i(\mathbf{X}, n_{Ai}) R_i(n_{Ai}) \; - \; C_A \qquad (1)$$

If attackers can direct their attacks against particularly vulnerable targets within the population of targets, then the fact that the expected profit from an attack in Equation 1 can differ across attackers raises potentially interesting questions about the additional incentives for defensive expenditure posed by the threat of such directed attacks. In order to focus on other incentive problems stemming from the dynamics of attacker choice, we make the following simplifying assumptions.

Suppose that attackers can observe the overall level of vulnerability for the population of targets but not the degree of vulnerability of any particular target. In this case, there is no reason for an attacker to attack one target instead of another. Or suppose, for some other reason, that each attacker directs attacks randomly against targets.

We further approximate a random assignment of attackers to targets by introducing the simplifying assumption that the total number of attacks, N_A, are spread uniformly over the N_T targets. In this case, $n_{Ai} = n_A = N_A/N_T$ for all i.

The cost C_A of mounting an attack includes the opportunity cost to the attacker of the lost profit from pursuing his or her next best choice. In this case, attackers should be motivated to launch attacks on the population of targets as long as the expected reward from launching an attack is greater than the cost of launching an attack. Hence, the equilibrium number of attacks per target, n_A^*, should satisfy the following equation:

$$\sum_{i=1}^{N_T} R_i(n_A^*)\, \sigma_i(\mathbf{X}, n_A^*)\, \frac{1}{N_T} \;=\; C_A \tag{2}$$

The left-hand side of Equation 2 represents the expected reward to an attacker from mounting an attack against the population of targets under consideration. The right-hand side of this equation is simply the cost to an attacker for mounting an attack. Equation 2 asserts that more attacks will be mounted until, in equilibrium, the expected reward from an attack equals the cost of the attack.

To simplify calculations, we suppose that a fractional number of attackers can choose to make attacks in any given period. In this case, the equilibrium number of attacks per target, n_A^*, will satisfy Equation 2 exactly rather than approximately as would be the case if n_A^* were required to be an integer.[6]

For Equation 2 to represent a reasonable model of attacker dynamics, the expected reward per attack on the left-hand side of Equation 2 should be a decreasing function of the number of attacks per target, n_A, at least near the equilibrium level n_A^*. That is,

$$\sum_{i=1}^{N_T} \left[\frac{dR_i}{dn_A}\, \sigma_i \;+\; R_i \frac{\partial \sigma_i}{\partial n_A} \right] \frac{1}{N_T} \;<\; 0 \tag{3}$$

for n_A close to n_A^*. For the “equal share” and “first winner takes all” models discussed earlier, $dR_i/dn_A \to -\infty$ as n_A becomes small, so that Equation 3 is satisfied for small enough values of n_A as long as $\partial \sigma_i/\partial n_A$ is bounded.

An important question that must be addressed in studies of attacker dynamics is: what determines the number of attacks? We believe that our model provides a promising answer to this question. The number of attacks is limited in our model by the competition for the rewards from a successful attack. This competition eventually reduces the expected reward to each attacker and thus provides a natural limit to the number of attacks.

Our explanation for what limits attacks will not be suitable in all circumstances. For example, in the “shark attack” model mentioned previously, the expected reward per attack against target i, $R_i(n_{Ai})\sigma_i(\mathbf{X}, n_{Ai})$ is an increasing rather than a decreasing function of n_{Ai}. In the shark attack model, there is no need for competition among

[6] To allow formally for a fractional number of attackers per target, we suppose that there are a continuum of potential attackers with the total mass of potential attackers equal to N_{PA}. The total number of actual attackers N_A must be less than or equal to N_{PA}.

attackers to divide the rewards from a successful attack. Moreover, the greater the number of attacks, the greater is the chance that at least one will be successful.

Note that the equilibrium level of attackers per target, n_A^*, satisfying Equation 2 depends in general on the vector of defensive expenditures by the N_T targets. The dependence of $n_A^*(\mathbf{X})$ on the levels of defensive expenditure chosen by the various targets has important implications for public policy.

2.3 Target Behaviour

Let L_i denote the expected value of the loss suffered by target i when one or more successful attacks on target i occurs. For simplicity, we suppose that L_i does not depend on the number of successful attacks but only on whether a successful attack occurs. This could be the case, for example, if the vulnerability which permitted a successful attack is patched after the first successful attack.

Targets are assumed to be risk neutral. Hence, we suppose that a target i will wish to choose its level of defensive expenditure x_i to minimize the expected loss

$$\sigma_i(\mathbf{X}, n_A^*)\, L_i \; + \; x_i \tag{4}$$

The objective in Equation 4 includes both the expected damage from an attack, $\sigma_i\, L_i$, and the cost of the defensive expenditure x_i, which the target must pay whether or not a successful attack occurs.

Note that the defensive expenditures of other targets can potentially affect target i's objective function in two ways. First, there may be ecosystem externalities. In addition, the equilibrium number of attacks against target i, n_A^*, typically depends on the entire vector of defensive expenditures $\mathbf{X}$. This dependence introduces a second type of externality through which target j's choice can affect target i.

2.4 Nash Equilibrium

For potential attackers, the expected payoff from an attack depends, in part, on the defensive choices of the targets. Similarly, the expected loss to a target depends on the number of attacks and, hence, on the choices of potential attackers. We model the strategic interaction between the choices of attackers and targets as a game.

A strategy for each potential attacker consists of a choice whether or not to participate in attacks on the population of targets. A strategy for target i is a choice of the level of defensive expenditure x_i. In a Nash equilibrium of the game between attackers and targets, the strategies of attackers and targets must be optimal given the expectations about the strategies held by other players. Moreover, these expectations must be correct when all parties behave optimally.[7]

[7] See, for example, Binmore (2007) or Myerson (1991).

Suppose that each player chooses its strategy simultaneously. When there are a large number of potential attackers, it will generally be optimal for some but not all such attackers to participate in attacks. Hence, each potential attacker must be indifferent between participating or not participating in an attack. For this to be the case, the equilibrium number of attacks per target in the Nash equilibrium, n_A^E, must satisfy the following equation:

$$n_A^E \; = \; n_A^*(\mathbf{X^E}) \tag{5}$$

where $n_A^*(\mathbf{X^E})$ on the right-hand side of Equation 17 is the solution to Equation 2 when the levels of defensive expenditure for the targets have been set equal to their Nash equilibrium levels which are denoted by $\mathbf{X^E}$.

Let $\mathbf{X_i}$ denote a vector of defensive expenditures by all targets except target i. Suppose that each target i chooses its level of defensive expenditure to minimize the expected loss in Equation 4 taking the number of attacks, n_A, and the levels of defensive expenditure of the other targets, $\mathbf{X_i}$, as fixed. For all values of n_A and $\mathbf{X_i}$, we suppose that the loss-minimizing level of defensive expenditure for target i is given by the usual first-order condition specified by the following equation:

$$-\frac{\partial \sigma_i(\mathbf{X}, n_A)}{\partial x_i} \, L_i \; = \; 1 \tag{6}$$

where the vector $\mathbf{X}$ in Equation 6 consists of the vector $\mathbf{X_i}$ of defensive expenditures for all targets except target i and the scalar x_i representing target i's defensive expenditure. Let $x_i(\mathbf{X_i}, n_A)$ denote the solution of Equation 6. We assume this solution exists and is unique for all $\mathbf{X_i}$ and n_A.

The Nash equilibrium levels of defensive expenditure, $\mathbf{X^E}$, satisfy the following equations for $i = 1, \ldots, N_T$:

$$x_i^E \; = \; x_i(\mathbf{X_i^E}, n_A^E) \tag{7}$$

where x_i^E is the equilibrium level of expenditure for target i, $\mathbf{X_i^E}$ denotes the equilibrium levels of defensive expenditure for all targets except i, and n_A^E is the equilibrium number of attacks per target. A Nash equilibrium of the game between attackers and targets is characterized by a quantity n_A^E and a vector $\mathbf{X^E}$ satisfying Equation 7, (in our worked example in Section 4 this will simplify to Equation 17).

3 A Role for Public Policy

In this section, we consider a policy-maker who wishes to minimize a weighted average of the expected losses suffered by the population of targets. Proposition 1 introduces the "incentive decomposition equation" that provides an interpretation of the benefits to society from additional defensive expenditure. Proposition 2 shows that the policy-maker will not generally wish to choose the level of expenditure determined in the unregulated Nash equilibrium.

Policy makers who may have the interest and ability to influence the defensive expenditures of targets include government regulators and law makers as well as the administrators of large computing facilities, such as cloud computing platforms, which are shared by the targets.[8]

In addition, once risk aversion has been introduced and insurance becomes an issue, insurance providers are also likely to have both the ability and desire to influence the level of defensive expenditure via the terms of insurance contracts offered to targets. Of course, the objectives of some of these parties are likely to be more complicated than simply the minimization of average losses.

Consider a policy-maker who can set the level of defensive expenditure for each target i and wishes to choose the vector of defensive expenditures for the targets to minimize the following weighted average of the targets' expected losses.

$$V \;=\; \sum_{i=1}^{N_T} v_i \left[\sigma_i(\mathbf{X}, n_A^*(\mathbf{X}))L_i \;+\; x_i\right] \tag{8}$$

where v_i are positive weights indicating how much importance the policy-maker places on the expected loss of target i. Without loss of generality, we can assume that

$$\sum_{i=1}^{N_T} v_i \;=\; 1\;. \tag{9}$$

Several features of the objective function in Equation 8 are worth noting. First, the formulation in Equation 8 implicitly assumes that target i bears the full cost of a successful attack on it. In principle, a successful attack on target i may impose losses on other members of society as well. For example, target i may provide services to consumers which are interrupted by a successful attack. One justification for our assumption is that target i chooses to provide compensation for all losses suffered by others from a successful attack. Target i might be motivated to provide such compensation because of liability laws or the desire to preserve a good reputation.

Uncompensated losses to other members of society from a successful attack on target i would comprise another type of technological or "ecosystem" externality. For simplicity, we do not consider this additional source of externalities in our analysis. Such externalities are straightforward to include and would not significantly affect the results of our analysis.[9]

Note that the equilibrium number of attacks per target is written as $n_A^*(\mathbf{X})$ in Equation 8. The equilibrium number of attacks per target is written in this way to emphasize that the policy-maker explicitly considers the effects of the targets' levels of defensive expenditure on attacker behaviour when choosing $\mathbf{X}$.

[8] See, for example, Motahari-Nezhad et al. (2009) and Pearson (2009), and for an overview from the insurance perspective see "Managing Digital Risk", Lloyd's 360 Risk Insight 2010, for a discussion of cloud computing and some of the risks associated with the cloud computing environment.

[9] Recall that the quantity L_i denotes the cost to target i from a successful attack. Let L_i^S denote the uncompensated costs borne by other members of society from a successful attack on target i. The effect of these losses can be included in our analysis by replacing L_i in Equation 8 by the sum $L_i + L_i^S$.

We suppose that the choice of **X** which minimizes the objective in Equation 8 satisfies the usual first-order conditions for an optimum. These first-order conditions can be rewritten as Equation 10 below. We refer to this equation as the "incentive decomposition equation" because it decomposes the marginal benefits from investment in defensive expenditure into four components. The extent to which a target is able to capture each component of benefit determines how the target's incentive to engage in such investment compares with that of the policy-maker.

Proposition 1: The incentive decomposition equation

$$-\frac{\partial \sigma_i}{\partial x_i} L_i \; + \; \left\{ -\frac{\partial \sigma_i}{\partial n_A}\frac{\partial n_A^*}{\partial x_i} L_i \right\} \; + \; \left\{ -\sum_{j\neq i}^{N_T} \frac{\nu_j}{\nu_i} \left\{ \frac{\partial \sigma_j}{\partial x_i} + \frac{\partial \sigma_j}{\partial n_A}\frac{\partial n_A^*}{\partial x_i} \right\} L_j \right\} = 1 \tag{10}$$

where, to simplify the notation, we have suppressed the dependence of the functions $n_A^*(\mathbf{X})$ and $\sigma_i(\mathbf{X}, n_A)$ on their various arguments.

Proof. The policy-maker's preferred choice of defensive expenditure for target i is obtained by setting the partial derivative of the objective function in Equation 10 with respect to x_i equal to zero. Dividing the resulting function by ν_i and rearranging terms produces the first-order condition for the optimal level of x_i that is specified in Equation 10. ■

Discussion of Proposition 1

Properties 1, 2, and 4 of the functions $\sigma_i(\mathbf{X}, n_{Ai})$ listed in Section 2 imply that the second and third the terms in on the left-hand side of Equation 10 are positive if the quantity $\partial n_A^* / \partial x_i$ is negative.

An expression for $\partial n_A^* / \partial x_i$ can be obtained by taking the partial derivative of the left-hand side of Equation 2 with respect to x_i and setting the result equal to zero.[10] Applying Properties 2 and 4 from Section 2 to the resulting equation shows that $\partial n_A^* / \partial x_i$ is negative if the inequality in Equation 3 is satisfied.

As in Equation 6, the first term on the left-hand side of Equation 10 represents the marginal reduction in expected damages to target i caused by the direct effect of an additional unit of defensive expenditure by target i on the probability σ_i. The direct effect occurs even when the number of attackers per target is fixed exogenously. As discussed in Section 2, $\partial \sigma_i / \partial x_i$ is assumed to be negative since an increase in the level of defensive expenditure should reduce the probability that at least one attack is successful. In this case, the first term is positive, which is what intuition would suggest.

[10] Since the function $n_A^*(\mathbf{X})$ satisfies Equation 2 for a range of **X**, the partial derivative with respect to x_i of the left-hand side of Equation 2 must equal the partial derivative of the right-hand side, which is zero.

The second term on the left-hand side of Equation 10 represents the marginal reduction in the expected damages to target *i* that occurs indirectly because an increase in target *i*'s defensive expenditure reduces the incentive for attackers to participate in attacks. This indirect benefit to target *i* is also positive.

In the Nash equilibrium described in Section 2.4, the assumption that players make their choices simultaneously implies that each target *i* takes the number of attacks per target as fixed when choosing its level of defensive expenditure. This assumption seems plausible since a change in any single target's defensive expenditure should have only a small effect on the overall vulnerability of the population of targets especially if the number of targets is large. In contrast, the policy maker sets the levels of defensive expenditures for all targets at once. Hence, it is plausible to assume, as we do, that the policy-maker takes into account the effect that changes in defensive expenditure would have on the behaviour of potential attackers.

Although, for simplicity, we considered a game in which targets and attackers made their choices simultaneously, we could have considered instead a game with sequential choices. In such a game, targets would first choose their levels of defensive expenditure and attackers would then make their choices about whether or not to participate in attacks. In such a game, a target would also consider the indirect benefit to itself which occurs through the effect its defensive expenditure might have on the equilibrium number of attacks per target. In this case, a term like that the second term in Equation 10 would also appear on the left-hand side of Equation 6. This change would not significantly affect our results. Whether individual choices are modeled with a game in which moves are simultaneous or sequential, an important role for public policy remains because of the final term on the left-hand side of Equation 10.

The expression within the final term on the left-hand side of Equation 10 represents the marginal reduction in the expected damages to targets other than target *i* resulting from an additional unit of defensive expenditure by target *i*. The terms $-\partial\sigma_j/\partial x_i$, for $j \neq i$ indicate the effects of possible ecosystem externalities which cause the defensive expenditure by target *i* to directly affect the probability that a successful attack is made on target j. Even when these terms are zero, the remaining terms in the summation on the left-hand side of Equation 10 indicate the benefit to other targets j caused because the defensive expenditure by target *i* reduces the overall profitability of attacks and, hence, the average number of attacks per target.

The final term on the left-hand side of Equation 10 create a diversion between the private benefits of defensive expenditure by target *i*, which are those obtained by target *i*, and the overall social benefits of this expenditure, which include the benefits to other targets. In the absence of policy intervention, an individual target is likely to have little or no incentive to take benefits for other targets into account when choosing its level of defensive expenditure.

Proposition 2: The role of public policy

The levels of defensive expenditure chosen by firms in the Nash equilibrium described in Section 2 are not socially optimal.

Proof. Proposition 2 is obtained by comparing Equation 6 and Equation 10. When the number of attackers per target is set equal to the level n_A^E which occurs in the Nash equilibrium, the Nash equilibrium levels of defensive expenditure, $\mathbf{X^E}$, satisfy Equation 6, for each i. Substituting $\mathbf{X^E}$ and n_A^E for $\mathbf{X}$ and $n_A(\mathbf{X})$ on the left-hand side of Equation 10 indicates that, at the Nash equilibrium levels of expenditure $\mathbf{X^E}$, the marginal reduction in the average expected damages produced by a small increase in x_i is greater than one. Hence, starting from the Nash equilibrium levels of expenditure and the Nash equilibrium number of attackers, the average expected loss can be reduced by increasing the expenditure of some target by a small amount. The levels of defensive expenditure in the Nash equilibrium cannot, therefore, be optimal from the policy-maker's point of view. ■

Corollary 2.1

If the number of attackers per target n_A is fixed exogenously and there are no ecosystem externalities, then the defensive expenditure of each firm in the Nash equilibrium is socially optimal.

Proof. When n_A is fixed exogenously and there are no ecosystem externalities (so that $\partial\sigma_j/\partial x_i = 0$ for all i and $j \neq i$), then the second and third terms in Equation 10 vanish. In this case, the first-order condition for the policy-maker's optimal choice for x_i reduces to Equation 6 which also describes the level of expenditure which target i would choose in the Nash equilibrium. ■

The policy-maker does not need to intervene to obtain his or her desired result when the number of attackers per target is fixed exogenously and there are no ecosystem externalities. The divergence between social and private benefits that motivates policy intervention occurs precisely because of the externalities introduced by the technological interactions that produce ecosystem externalities and by the response of attackers to the overall vulnerability of the population of targets.

4 A Policy Example with a First-Winner-Takes-All Model and Identical Targets

This section presents a concrete example that illustrates the main ideas of the paper. In addition to providing insight into the conclusions of our paper, the analysis in this section can be regarded as providing an alternate justification for these conclusions.

In order to focus on the externalities caused by the behaviour of attackers, we also assume in this example that there are no ecosystem externalities. Hence, the probability that one or more attacks on target i are successful depends only on the defensive expenditure of target i and the number of attacks on target i, n_{Ai}, and we write this probability as $\sigma_i(x_i, n_{Ai})$.

Moreover, we further assume that targets are identical before the levels of defensive expenditure are chosen. Hence we can suppress the subscript i and write the above function as $\sigma(x_i, n_{Ai})$ for all i. We assume that the function $\sigma(x_i, n_{Ai})$ has the following functional form,

$$\sigma(x_i, n_{a_i}) \;=\; 1 - \exp(\ln(1-\alpha(x_i))n_{Ai}) \tag{11}$$

where

$$\alpha(x_i) \;=\; A_0\, e^{-a_0 x_i} \tag{12}$$

a_0 is a positive constant and $A_0 \in (0, 1]$.

The motivation for Equation 11 is as follows. Suppose that $\alpha(x_i)$ denotes the probability that a single attack on target i will be successful. In this case, the quantity A_0 represents the probability that a single attack on target i is successful in the absence of defensive expenditure, i.e., $x_i = 0$, and a_0 in represents the level of investment by a target that reduces the probability that a single attack is successful by the factor $1/e$. If the outcomes of individual attacks represent independent events and there are a total of n_{Ai} attacks on target i, then $\sigma(x_i, n_{Ai})$ must satisfy the following equation:

$$\sigma(x_i, n_{Ai}) \;=\; 1 - (1-\alpha(x_i))^{n_{Ai}} \tag{13}$$

Equation 11 is simply the interpolation of Equation 13 to non-integer values of n_{Ai}.[11]

Targets are assumed to be risk neutral. As in Section 2, we model the unregulated behaviour of attackers in terms of a game between attackers and targets where all players move simultaneously.

Since targets are assumed to be identical, the loss incurred due to a successful attack, L, is assumed to be the same for all targets. Target i chooses its level of defensive expenditure, x_i to minimize the expected loss, $\sigma(x_i, n_{Ai})L + x_i$ taking the number of attacks n_{Ai} as fixed. Let $x^*(n_{Ai})$ denote the solution to this minimization problem for a given n_{Ai}. For the relevant range of values for n_{Ai}, this solution exists and is unique for the parameter values and functional forms chosen in our example. We can write the following equation:

$$x^*(n_{Ai}) = \arg\min_x \left(\sigma\left(x, n_{Ai}\right) L + x\right) \tag{14}$$

Note that, since a target's expected loss does not depend on i, the function $x^*(n_{Ai})$ is the same for all targets i. We will refer to this function as a target's "best reply function" or "best response function" since it represents an optimal response (from an individual target's point of view) to a level of attack n_A. Although, in principle, different targets are free to choose different levels of defensive expenditure, it is

[11] As a check on our analysis, we will sometimes find it convenient to also consider the following second order Taylor approximation for Equation 11. $\sigma(x_i, n_{Ai}) \;=\; n_{Ai}\left[\alpha(x_i) + \frac{(\alpha(x_i))^2}{2}\right] - n_{Ai}^2 \frac{(\alpha(x_i))^2}{2}$.

optimal for identical targets to choose the same level of defensive expenditure when faced with the same level of threat, n_A.

As in Section 2, we suppose that each potential attacker makes at most one attack and there is an equal chance that an attacker attacks any given target. Moreover, we assume that the expected reward from a successful attack is described by a first-winner-takes-all model. Suppose that each target chooses the same level of defensive expenditure, x, and the number of attacks per target is n_A. In this case,the expected profit obtained by an attacker who chooses to mount an attack is given by the objective function

$$\sigma(x,n_A)\,\frac{R}{n_A}\;-\;C_A \qquad (15)$$

where R denotes the reward obtained by an attacker who makes the first successful attack on a target, $\sigma(X,n_A)/n_A$ represents the probability that an attacker is the first of the n_A attackers to make a successful attack, and C_A is the cost to an attacker from making an attack. We continue to assume, for simplicity, that potential attackers are identical so that C_A is the same for all attackers. Equation 15 is the special case of Equation 2 in Section 2 which is relevant for our example.

As discussed in Section 2, a strategy for each potential attacker in the Nash equilibrium of the game we consider is a choice of whether or not to mount an attack. With a large number of potential attackers, the two choices must be equally profitable so that some but not all potential attackers choose to participate in attacks. Hence, when each target chooses the same level of defensive expenditure, x, the equilibrium number of attackers per target, $n_A^*(x)$, must satisfy the following equation:

$$\frac{\sigma(x,n_A^*)}{n_A^*}\;=\;\gamma \qquad (16)$$

where $\gamma = C_A/R$ is the cost of an attack per unit of reward. Equation 16 is obtained by setting the objective in Equation 15 equal to zero and dividing by R. We refer to $n_A^*(x)$ as the attacker's response function since it describes how attackers respond to the level of defensive expenditure chosen by the targets.[12]

Since all targets face the same number of attacks, all targets will choose the same level of defensive expenditure in the Nash equilibrium. The Nash equilibrium is therefore described by a pair of numbers, x^E and n_A^E satisfying the following equations.

$$x^E = x^*(n_A^E)\;\text{ and }\;n_A^E = n_A^*(x^E) \qquad (17)$$

[12] For the quadratic approximation for σ, it is straightforward to compute an analytic formula for the attacker response function. Simply substitute the approximate formula for σ into Equation 16 and solve the resulting equation to obtain the following formula for $n_A^*(x)$:

$$n_A^*(x) \;=\; \frac{2\,(\alpha(x)-\gamma)}{(\alpha(x))^2}\;+\;1$$

The results for this approximation are plotted as dashed lines in the graphs used to illustrate our example.

Equation 17 asserts that the defensive expenditure of each target, x^E, is optimal given the targets' beliefs about the level of attacks. In addition, the choices of potential attackers are optimal given their beliefs about the vulnerability of the population of targets. Moreover, the beliefs of the targets and potential attackers are consistent with the actual choices made in the Nash equilibrium.

Graphically, the equilibrium values, x^E and n_A^E occur at the intersection of the two curves, $x^*(n_A)$ and $n_A^*(x)$.

4.1 Baseline Model Parameters

The version of the model in our example involves only four exogenous parameters A_0, a_0, L and $\gamma = C_A/R$. For the remainder of this section our baseline case is that the probability of a single successful attack given no defensive expenditure is $A_0 = 0.1$, the $1/e$-life effectiveness of extra defensive expenditure is $a_0 = 0.5$, the loss from an attack is $L = 10$ units and the cost of attack to reward ratio is $\gamma = 0.01$.

Figure 1 graphs the probability σ that one or more successful attacks are made on a given target for the parameter values specified above. Each curve shows how this probability declines as the defensive expenditure x increases for selected numbers of attackers (each executing a single attack) n_A.

Figure 2 presents the expected loss, $\sigma L + x$, for a target as a function of the defensive expenditure x. Each curve shows the expected loss for a different value of n_A.

For each indicated value of n_A, the optimal level of defensive expenditure for a target, $x^*(n_A)$, occurs where the curves in Figure 2 take on their minimum values. From Figure 2 we observe that the shape of the expected loss function is well behaved and the optimal level of defensive expenditure tends to zero as the number of attackers becomes small.

Figure 3 shows the expected profit per unit of reward for an attacker, $(\sigma(x, n_A)/n_A)$ $- \gamma$, for the baseline scenario as a function of the total number of attackers. Each curve shows the expected profit for a given level of defensive expenditure as a function of the number of attackers. The equilibrium level of attacks per target, $n_A^*(x)$, occurs where the expected profit is zero so that potential attackers are indifferent between participating and not participating in attacks.

The two response curves $n_A^*(x)$ and $x^*(n_A)$, are plotted in Figure 4. The Nash equilibrium values, x^E and n_A^E, occur at the intersection of these two curves. Note that the intersection of the dashed lines representing the quadratic approximation is very close to the intersection of the solid lines.

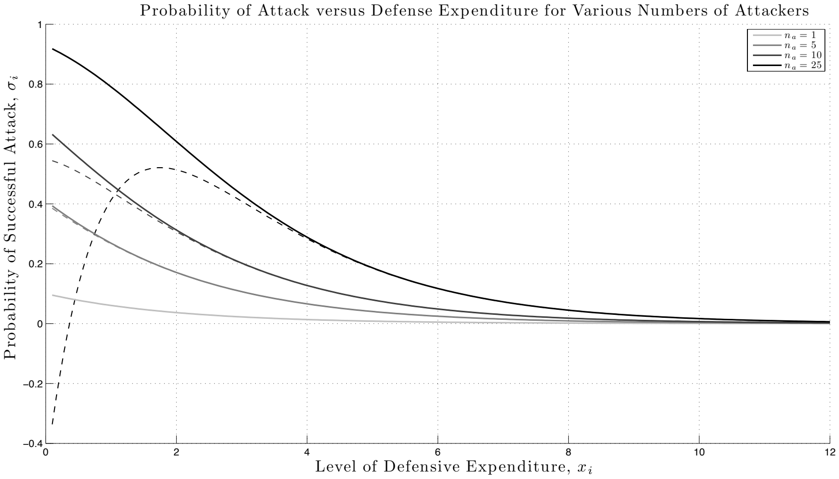

Fig. 1 The probability of a successful attack on a single target as a function of defensive expenditure x for several different numbers of attackers n_A. The unbroken line plots the interpolated function from Equation 11. The dashed line presents the function using a quadratic approximation. For this graph, $A_0 = 0.1$ and $a_0 = 0.5$. These values will be used throughout the example.

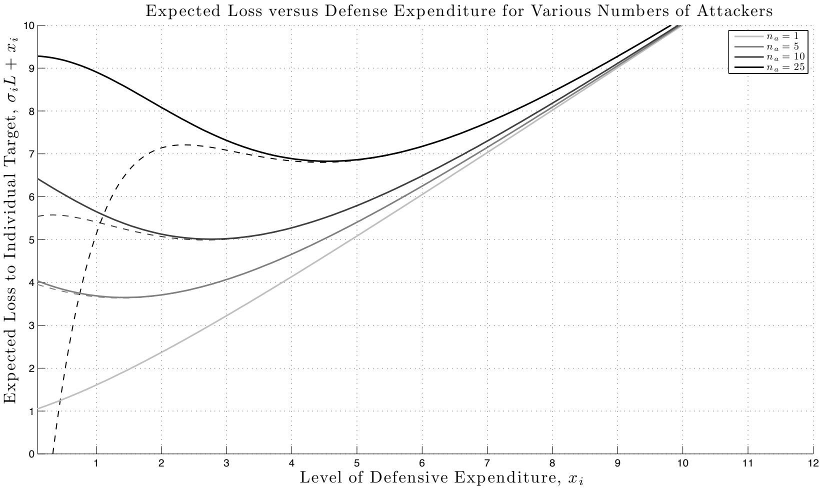

Fig. 2 The expected loss of a target, $\sigma_i L + x_i$, versus defensive expenditure of the target for various numbers of attackers. The dashed line represents the expected loss calculated using the quadratic approximation of σ and the unbroken line uses the interpolated version of σ. For this graph, we set the value of the loss from a successful attack to $L = 10$ and this value will be used throughout the example.

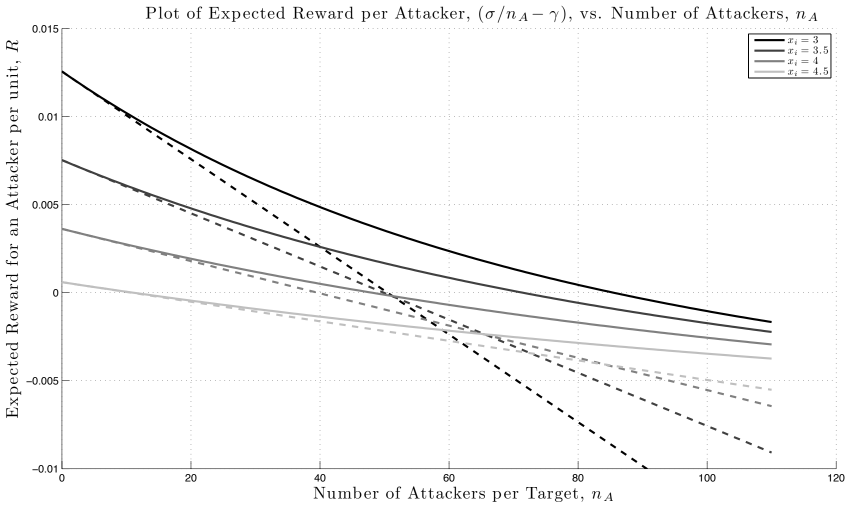

Fig. 3 The expected reward per attacker versus the number of attackers for selected levels of defensive expenditure by the targets. The cost reward ratio $\gamma = 0.01$ is used for this example. For a given level of defensive expenditure the equilibrium number of attackers per target occurs where the expected reward curve intersects a horizontal line through zero.

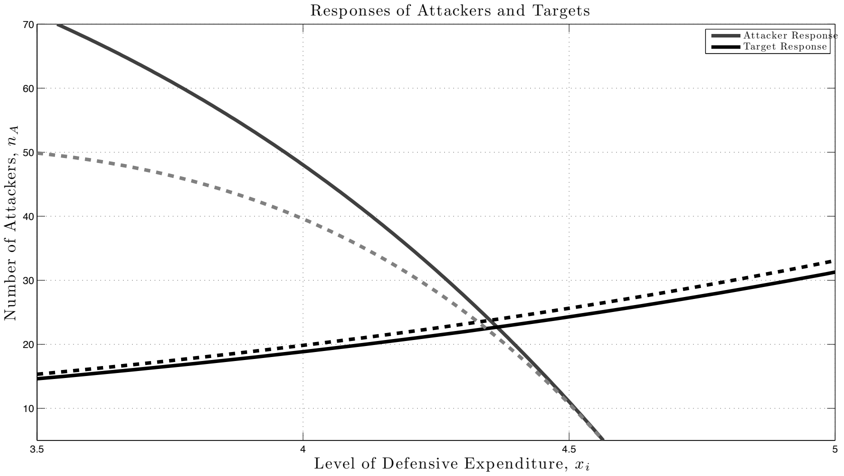

Fig. 4 The response curves for attackers and targets. The dashed and unbroken lines represent, respectively, calculations using the quadratic and interpolated versions of the σ function. The assumed parameters for this example are $L = 10$, $\gamma = 0.01$, $A_0 = 0.1$ and $a_0 = 0.5$. The Nash equilibrium occurs at the point of intersection of the attacker and target response curves.

4.2 The Policy-maker's Problem

We now consider a benevolent policy-maker who wishes to minimize the expected loss per target. In Section 3, we considered a general weighted average. Since the targets in our example are identical, we restrict attention to the case where the expected loss for each target is given the same weight. In the notation of Section 3, $\nu_i = 1/N_T$ for all i.

For ease of comparison with the Nash equilibrium, we also restrict attention to the case where the policy maker specifies the same level of defensive expenditure for each target. If the level of defensive expenditure, x^P, which the policy maker finds optimal to specify for each target differs from the Nash equilibrium level of defensive expenditure, x^E, that is sufficient to establish a role for compulsory regulation. In what follows, we show that $x^P > x^E$ for our example.

Substituting $\nu_i = 1/N_T$ and $x_i = x$ into Equation 8 produces the following equation for the policy-maker's objective function.

$$V(x) = \sigma(x, n_A^*(x))L + x \tag{18}$$

Equation 18 indicates that the policy-maker wishes to choose a level of defensive expenditure x to minimize the expected loss of a typical target taking into account the effect that a change in defensive expenditure has on the behaviour of potential attackers.

Superficially, the policy-maker's objective in Equation 18 appears to resemble closely the objective of an individual target in Equation 14. Hence, one might be led to speculate that any divergence between the policy-maker's choice, x^P, and the Nash equilibrium value, x^E, is not due to a divergence between private and social incentives, as we have asserted in Section 3, but is simply an artifact of the assumption that all players make their choices simultaneously in the Nash equilibrium.

To understand the difference between the objective in Equation 18 and the objective in Equation 14, it is important to note that the function $n_A^*(x)$ indicates the equilibrium level of attackers per target when the level of defensive expenditure for each and every target is set equal to x. In other words, the vector of defensive expenditures is the constant vector $\mathbf{X} = x, x, \ldots, x$, for all $i = 1, \ldots, N_T$. In contemplating a change from the level x to x', the policy-maker is contemplating a change not in the level of defensive expenditure for one target but a change in the entire vector of defensive expenditures from $\mathbf{X}$ to $\mathbf{X}' = x', x', \ldots, x'$.

To describe properly a Nash equilibrium where targets move first and, therefore, take into account the effect of their individual expenditures on the behaviour of attackers, a new function $n_A^*(x_i, \mathbf{X_i})$ would have to be defined. This function would indicate how the equilibrium number of attackers changes when the level of defensive expenditure for target i, x_i, varies while the levels of defensive expenditure for all other targets other than i, $\mathbf{X_i}$, are held fixed. Substituting $n_A^*(x_i, \mathbf{X_i})$ for n_{Ai} on the right-hand side of Equation 14 would produce the objective which is minimized by a target in a game with sequential moves. This objective is by no means the same as the objective in Equation 18. Moreover, as demonstrated in Section 3, the rationale for regulation would not be significantly affected if the unregulated choices of targets were described by a Nash equilibrium where targets moved first.

There is one special case in which the objectives in Equation 18 and Equation 14 are in fact the same. This is the case where attackers do not respond to the vulnerability of the population of targets so that the number of attackers per target is fixed exogenously at some level n_A. In this case, $n_{Ai} = n_A$ in Equation 14 and $n_A(x) = n_A$ in Equation 18. For this special case, the optimization faced by the policy-maker is the same as the optimization faced by an individual target. As was also observed in Section 3, when there are no technological externalities and the number of attackers is fixed exogenously, the equilibrium levels of defensive expenditure chosen by targets is the same as those which would be chosen by the policy-maker, and no policy intervention is required.

In all other cases, the policy-maker must take into account the response of attackers when choosing which level of x to specify for the targets. In particular, the optimal choice for the policy-maker, x^P, satisfies the following equation:

$$x^P = \arg\min_x \sigma(x, n_A^*(x))L + x \tag{19}$$

The dashed and solid curves in Figure 5 show the policy-maker's loss function, $V(x)$, with and without the quadratic approximation for $\sigma(x, n_A)$. The square boxes, where the two curves start, represent the predicted outcomes in the Nash equilibrium.

As shown in Figure 5, it turns out that for our example the policy-maker's loss function is declining throughout the range of defensive expenditures starting at the Nash equilibrium levels and ending at a level just below that which deters all attacks. Hence, the optimal choice of defensive expenditure for the policy-maker is certainly not x^E. By increasing the level of defensive expenditure from the Nash equilibrium level to the highest level in this range, the policy-maker can reduce the expected loss for each target from a level above 6.5 to a level below 5, a reduction of approximately 25%.

In this example, the policy-maker can substantially reduce the average loss suffered by each target by prescribing higher levels of defensive expenditure than those which would be chosen at the Nash equilibrium. Indeed, for this example, the policy-maker would find it optimal to choose a mandatory level of defensive expenditure that would virtually eliminate the threat from attackers.

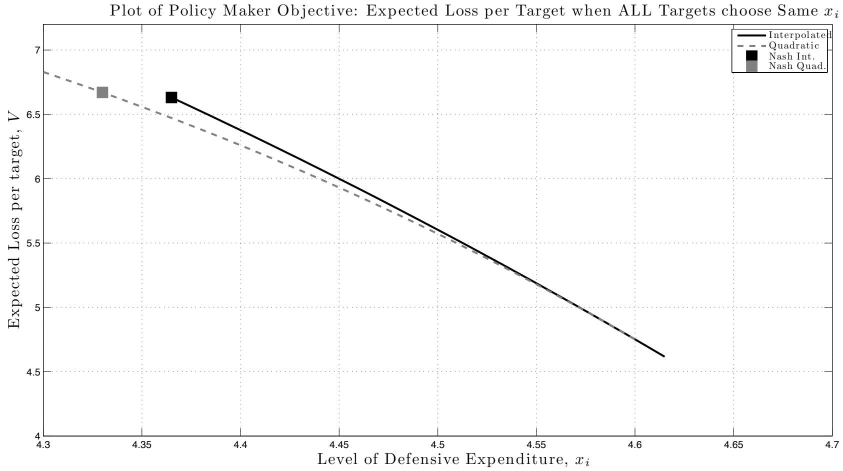

Fig. 5 The policy-maker's loss function, $V(x)$. The curves are truncated below the point where $n_A = 1$. Since the curves are continuously decreasing for the relevant range of expenditures it is optimal for the policy-maker to prescribe a value for x corresponding to the lowest point on the curve. By doing so, the policy-maker reduces the expected loss to each target by approximately 25% compared to the Nash equilibrium outcomes which are indicated by the square markers.

5 Concluding Remarks

Our paper offers a contribution to the ongoing debate on the appropriate form of public policy with respect to cyber threats. The paper introduces a model in which attackers and targets interact strategically. In this model, the levels of defensive expenditure chosen by targets acting independently are not socially optimal. This suggests a rationale for regulations that set compulsory minimum levels of investment in security.

The divergence between private and public incentives that provides the rationale for regulation occurs because investments in security by one target provide "external" benefits for other targets. Our paper identifies two channels by which such externalities can occur.

First, there are technological externalities which occur because a vulnerability in the software or hardware operated by one target may facilitate attacks on other targets. In addition, the losses from a successful attack may not be incurred only by the target that has suffered the attack.

A second type of externality occurs because of the strategic interaction between targets and attackers. In particular, an investment in security by one target is likely to make attacks on the entire population of targets less attractive. The worked example in Section 4 of the paper focuses on this externality and illustrates with a plausible set of assumptions that public policy interventions can produce substantial reductions in the average loss across all targets.

References

Baldwin, A., I. Gheyas, C. Ioannidis, D. Pym, and J. Williams (2012). Contagion in cybersecurity attacks. In R. Böhme (Ed.), *Workshop on the Economics of Information Security 2012*. WEIS.

Binmore, K. (2007). *Playing for Real; A Text on Game Thoery*. Oxford University Press.

Cavusoglu, H., H. Cavusoglu, and J. Zhang (2008). Security patch management: Share the burden or share the damage. *Management Science 54*(4), 657–670.

Chen, P., G. Kataria, and R. Krishnan (2011). Correlated failures, diversification, and information security risk management. *Management Information Systems Quarterly 35*(2), 397–422.

Cornes, R. and T. Sandler (1996). *Theory of Externalities, Public Goods, and Club Goods* (2nd ed.). Cambridge University Press.

Cremonini, M. and D. Nizovtsev (2010). Risks and benefits of signalling information system characteristics to strategic attackers. *Journal of Management Information Systems 26*(3), 241–274.

Florencio, I. and C. Herley (2011). Where do all the attacks go? In B. Schneier (Ed.), *Workshop on the Economics of Information Security 2011*. WEIS.

Fultz, N. and J. Grossklags (2009). Blue versus red: Towards a model of distributed security attacks. In R. Dingledine and P. Golle (Eds.), *Proc. Financial Cryptography and Data Security '09*, Volume 5628 of *LNCS*, pp. 167–183. Springer.

Gordon, L. and M. Loeb (2002). The economics of information security investment. *ACM Transactions on Information and Systems Security 5*(4), 438–457.

Gordon, L. A., M. P. Loeb, and T. Sohail (2010). Market value of voluntary disclosures concerning information security. *MIS Quarterly 34*(3), 567–594.

Herley, C. (2012). Why do Nigerian Scammers say they are from Nigeria? In R. Böhme (Ed.), *Workshop on the Economics of Information Security 2012*. WEIS.

Ioannidis, C., D. Pym, and J. Williams (2009). Investments and trade-offs in the economics of information security. In R. Dingledine and P. Golle (Eds.), *Proc. Financial Cryptography and Data Security '09*, Volume 5628 of *LNCS*, pp. 148–166. Springer. Preprint available at `http://homepages.abdn.ac.uk/d.j.pym/pages/IoannidisPymWilliams-FC09.pdf`.

Ioannidis, C., D. Pym, and J. Williams (2011). Information security trade-offs and optimal patching policies. *European Journal of Operational Research 216*(2), 434–444.

Ioannidis, C., D. Pym, and J. Williams (2012). Fixed costs, investment rigidities, and risk aversion in information security: A utility-theoretic approach. In B. Schneier (Ed.), *Economics of Security and Privacy III*. Springer. Proceedings of the 2011 Workshop on the Economics of Information Security.

Laffont, J. J. (2008). *Fundamentals of Public Economics*. MIT Press Books.

Motahari-Nezhad, H., B. Stephenson, and S. Singhal (2009). Outsourcing business to cloud computing services: Opportunities and challenges. Technical report, Hewlett Packard Laboratories Working Paper HPL-2009-23.

Myerson, R. (1991). *Game Theory: Analysis of Conflict*. Harvard University Press.

Pearson, S. (2009). Taking account of privacy when designing cloud computing services. Technical report, Hewlett Packard Laboratories Working Paper HPL-2009-54.

Straub, D. W. and R. J. Welke (1998). Coping with systems risk: Security planning models for management decision making. *Management Information Systems Quarterly 22*(4), 441–469.

Varian (2004). System reliability and free riding. Available at: `http://people.ischool.berkeley.edu/~hal/people/hal/papers.html`.

Varian (2010). *Intermediate Microeconomics: A Modern Approach*. WW Norton & Co.