SECURITY IN PARTNERSHIPS[*]

Michael H. Riordan
Department of Economics, Columbia University

October 2014

I. Introduction

In 2013, Target Corporation's business relationship with an HVAC company compromised the confidential information of more than 110 million consumers. By first infecting the computer system of the outsourced HVAC vendor, cyber attackers stole digital credentials, enabling them to intrude into Target's data systems and to steal credit card data that Target had acquired from consumers. The apparently weak security precautions of the HVAC company exposed Target to an indirect attack against which its own security precautions were lacking (U.S. Senate Committee on Commerce, Science, and Transportation, 2014).

Partners in different kinds of relationships confront interdependent security risks with a similar structure. While the Target breach involved consumer data theft, other cybersecurity threats involve economic espionage, intellectual property theft, or sabotage.[1] Interdependent airline security is another case in point, because baggage and

---

[*] This research is supported by the Digital Society Initiative of the Institute for Advanced Study in Toulouse (IAST) for which I was a Fellow in 2013-14. Thanks go to Rishab Guha and Ilton Soares for superb research assistance.

[1] For a prominent example of economic espionage, see "Chinese Army Unit is Seen as Tied to Hacking Against U.S.," *New York* Times, February 18, 2013, and "5 in Chinese Army Face U.S. Charges of Cyberattacks," *New York Times*, May 19, 2014. For an example of an indirect attack, see Mandiant (2013)'s case study of an energy company breached by an attacker stealing proprietary data. The attacker is surmised to have gained access to the energy company's internal network initially through an outsourced IT service provider whom the attacker had breached already. After remediation of this

passengers possibly carrying a bomb transfer from one aircraft to another for connecting flights (Kunruether and Heal, 2003; Heal and Kunreuther, 2007). A mundane example is a married couple concerned about the flu. If one catches the flu, the other is more likely to get it. The similarity of these quite different scenarios is that, if an external threat breaches the defenses of one partner, then the other partner is exposed to a heightened risk of breach.

Interdependent security poses an organizational problem. How can a partnership address externalities in the decentralized provision of security? The externality problem is easy to understand, even though its solution may be elusive. A positive externality arises if security precautions taken by one party also increase the security of its partner. Partners who incur the costs of their own precautions while sharing the benefits can be expected to underinvest in security unless there is a mechanism to align incentives. For husband and wife, incentive alignment might come from devotion. For firms in a business relationship, a contract to share the losses or to reward good outcomes might be the answer.

This paper examines the interdependent security problem in a game-theoretic model of a partnership. Each partner has an expected loss from a security breach, and each invests in precautions to reduce its vulnerability. There are two distinct risks. A breach can result either from a direct attack by an external threat, or from an indirect attack channeled through its compromised partner. Security is interdependent because a firm's risk of an indirect attack depends on its partner's security against direct attacks. Thus, for better or worse, the partners share responsibility for each other's security.

---

initial security breach, a second attack was launched through a third company with whom the energy company had a business partnership. The Stuxnet worm, which disabled Iran's nuclear centrifuges, is a famous example of cybersabotage; see "Obama Order Sped Up Wave of Cyberattacks against Iran," *New York Times,* June 1, 2012. Denial of service attacks which threaten to disable websites are common; see, for example, "Distributed Denial of Service," *New York Times,* April 1, 2013.

The partners' incentives to invest in security might depend on the economies of scope in reducing the risks of direct and indirect attacks. Previous economics literature on the subject has focused on two extreme cases. In an "exogenous contagion model" the probability of an indirect security breach from a compromised partner is exogenous. Alternatively, in a "balanced security model", costly precautions protect equally well against direct and indirect attacks. More generally, security investments might reduce asymmetrically the risks of direct and indirect attacks. The model studied here generalizes in a straightforward way by treating security as a composite good. In other words, security from indirect attacks is assumed to be proportional to security from direct attack. The proportionality factor indicates the relative difficulty of protecting against indirect versus direct attacks. The exogenous contagion and balanced-security models are extreme cases of the composite-security model.

The rest of this paper is organized as follows. Section II introduces a conceptual framework and selectively reviews previous literature on security economics. Section III introduces the composite-security model. Assuming diminishing returns to security precautions, Section IV characterizes best responses and Section V establishes the existence of a unique equilibrium. Assuming a well-behaved social welfare function, Section VI characterizes optimal security and Section VII studies how penalties and bonuses can solve the externality problem. Section VIII discusses standards in the context of cybersecurity, and Section IX concludes with directions for future research.

II. General Framework and Literature Review

Consider two firms in a business partnership. For simplicity, suppose that both face symmetric security threats. Let α denote the independent probability of a direct attack against each firm, and β the probability of an indirect attack through a compromised partner. Define security as the probability that a firm is not breached, and let x_i and y_i respectively denote the levels of Firm i's precaution against direct and indirect attacks. Then the total security for Firm i is equal to the joint probability that it is safe from both

kinds of attacks, which depends on its own precautions against direct and indirect attacks and on its partner's (Firm j's) precautions against direct attacks:

$$s(x_i, y_i, x_j) = [1-\alpha(1-x_i)][1-\alpha\beta(1-x_j)(1-y_i)]$$

Security is interdependent because a firm's total security depends on its partner's security against direct attacks.

The objective of each firm is to maximize its utility of security. Let $\theta_i \geq 0$ denote Firm i's (expected) loss from a security breach. For simplicity, assume that each firm has the same cost function $c(x_i, y_i)$ for achieving security from direct and indirect attacks. Then the utility of precautions for Firm i depends on its total security and the cost of its own precautions:

$$u_i = \theta_i s(x_i, y_i, x_j) - c(x_i, y_i)$$

Because security is interdependent, the marginal utility of x_i is increasing in x_j, while the marginal utility of y_i is decreasing in x_j. As we shall see, the opposite signs of these cross partial derivatives are responsible for some ambiguity about whether investments in security are strategic complements or strategic substitutes (Bulow, Geanakoplos, and Klemperer, 1985).

This general framework can be interpreted to embed Gordon and Loeb (2002)'s decision-theoretic analysis of security investment in a two-person game. If there were no risk of contagion ($\beta = 0$), then the security of Firm i would depend exclusively on its own investment against direct attacks:

$$s(x_i) = 1-\alpha(1-x_i)$$

Since there is no need to defend against indirect attacks, the cost of security depends only on precautions against direct attacks:

$$c(x_i) \equiv c(x_i, 0)$$

Gordon and Loeb (2002) argue that the optimal investment in security is not necessarily monotonic in baseline "vulnerability". Suppose there is a baseline level of security $s(\chi)$ corresponding to zero investment, i.e. $c(\chi) = 0$. Gordon and Loeb

(2002) define vulnerability to be equal to $1-\chi$, and study how shifts in the cost function that increase vulnerability can alter the incentive for security investment. For the quadratic cost function introduced later in this paper, however, security investment is always increasing in χ.

Varian (2004) views interdependent security as a privately-provided public good. Firm i contributes to the public good by investing in precautions x_i, and both partners enjoy the same security level $s(x_i, x_j)$. Building on Hirschliefer (1983), Varian (2004) considers three different assumptions about the security function: (1) in the "aggregate effort" model, $s(x_i, x_j) = x_i + x_j$; (2) in the "best-shot" model, $s(x_i, x_j) = \max\{x_i, x_j\}$; and (3) in the "weakest link" model, $s(x_i, x_j) = \min\{x_i, x_j\}$. The cost of precautions is assumed to be a convex increasing function $c(x_i)$.[2] In a complete information Nash equilibrium, Firm i chooses a best response:

$$x_i = \arg\max\{\theta_i s(x, x_j) - c(x) \mid 0 \le x \le 1\}$$

Varian (2004) derives how equilibrium allocations of responsibility vary with the security function. In the aggregate-effort and best-shot models, only the high-value partner invests, and imposing on this agent a fine equal to his partner's losses achieves a social optimum.[3] In the weakest-link model there are Pareto-ranked multiple equilibria in which the partners invest equally, and the Pareto-superior equilibrium is determined by the firm with the least incentive for security.

Kunreuther and Heal (2003) treat interdependent security as a coordination problem. Each firm can make a discrete costly investment to eliminate completely the risk from a direct attack, but remains defenseless against an indirect attack.[4] This is equivalent to fixing $y_i = 0$ and restricting $x_i \in \{\chi, 1\}$ with $c(\chi) = 0$ and $c(1) > 0$. The parameter β now is interpreted as an exogenous contagion risk. In

---

[2] Varian (2004) also allows for asymmetric investment costs.
[3] In the case of asymmetric investment costs, the optimal fine should be imposed on the low-cost partner (Varian, 2004).
[4] See also Heal and Kunreuther (2007).

contrast to Varian (2004), interdependent security is not a pure public good. Instead, the security of each partner might depend asymmetrically on its own investment decision (x_i) on its partner's decision (x_j):

$$s(x_i, x_j) = [1 - \alpha(1 - x_i)][1 - \alpha\beta(1 - x_j)]$$

This function is supermodular, i.e. the cross partial derivative is positive. Consequently, security investments are strategic complements, and, for a symmetric environment ($\theta_i = \theta_j$), either both firms invest in security ($x_i = x_j = 1$) in equilibrium, or neither do ($x_i = x_j = \chi$). Furthermore, multiple equilibria are possible.

Acemoglu, Malekian, and Ozdaglar (2013) view interdependent security as network security problem, under the assumption that security investments protect equally well against direct and indirect attacks. In the case of a simple partnership, this balanced security assumption amounts to setting $y_i = x_i$ and letting $c(x_i) \equiv c(x_i, x_i)$. Consequently, the security function simplifies to: [5]

$$s(x_i, x_j) = [1 - \alpha(1 - x_i)][1 - \alpha\beta(1 - x_i)(1 - x_j)]$$

This function is not globally supermodular, i.e. the cross-partial derivative is negative or positive, depending on whether x_i is greater or less than $\frac{1}{2}$. Therefore, as we shall see more clearly later, security investments might be either strategic complements or strategic substitutes in equilibrium.

Acemoglu, Malekian, and Ozdaglar (2013)'s concern, however, is not with a simple two-firm partnership, but rather with interdependent security in multi-agent random networks. The paper points out that a general presumption of underinvestment in security in this context is unwarranted if agents are positioned asymmetrically in the network and investments are strategic substitutes. Furthermore, it is possible in equilibrium for an

---

[5] Acemoglu, Malekian, and Ozdaglar (2013) assume mutually exclusive risks of direct attacks rather than independent risks. Consequently, in total security Acemoglu, Malekian, and Ozdaglar (2013) is additive in security from direct and indirect attacks, rather than multiplicative as represented here.

agent to "overinvest" in security relative to the social optimum, in order to compensate for the deficient investment of other agents to whom it is connected, and, consequently, equilibrium might be "more secure" than the welfare optimum.

The distinction between security against direct and indirect attacks raises questions about the technology for providing security. Is it easier to guard against direct attacks than indirect attacks? Are there economies of scope in provision of the two kinds of security? The above selective survey of the security economics literature suggests that incentives for providing security are sensitive to the technology for precautions against attacks. In the exogenous contagion model of Kunreuther and Heal (2003), security investments are strategic complements, whereas in the balanced security model of Acemoglu, Malekian, and Ozdaglar (2013) investments "typically" are strategic substitutes. A generalization of the exogenous contagion and balanced security models, pursued in the next section, assumes an arbitrary proportional relationship between protections from the two kinds of attacks. This composite-security model simplifies the general framework by supposing a particularly strong form of economies of scope that essentially makes security investment a one-dimensional choice problem for which standard scale economy concepts are well-defined (Baumol, Panzar, and Willig, 1982). The exposition is simplified further by maintaining assumptions guaranteeing that relevant objective functions are sufficiently well behaved so that optima are determined by first-order conditions.

III. Composite Security Model

Security now is assumed to be a composite good, providing security against direct and indirect attacks in a constant proportion:

$$y_i = \phi x_i$$

Given $0 \leq \phi \leq 1$, the total security of Firm i depends both on its own and its partner's investments in composite security:

$$s(x_i, x_j) = [1 - \alpha(1 - x_i)][1 - \alpha\beta(1 - x_j)(1 - \phi x_i)]$$

Exogenous contagion is the special case $\phi = 0$, and balanced security is $\phi = 1$. Intermediate cases indicate varying degrees of greater difficulty to protect against indirect attacks compared to direct attacks.

Firm i suffers a loss θ_i from a breach resulting from either a direct or indirect attack. In line with previous literature, the magnitudes of these losses are assumed to be common knowledge. For simplicity, the cost of security is assumed to be quadratic:

$$c(x_i) = \tfrac{1}{2}(x_i - \chi)^2$$

The parameter χ is the baseline security if the firm does not invest in precautions. The utility of security precautions is the expected benefit of preventing a breach less the cost of precautions:

$$u(\theta_i, x_i, x_j) \equiv \theta_i s(x_i, x_j) - c(x_i)$$

The simplifying advantage of the composite-security model is that investments in security for each partner are one-dimensional, i.e. x_i indexes security precautions against both direct and indirect attacks, and ϕ indicates the relative effectiveness of these precautions against indirect attacks. The linear-quadratic specification allows convenient closed-form solutions.

Diminishing returns to investments in security means that an additional dollar spent on precautions is less beneficial the more is spent, i.e. the marginal utility for Firm i is declining in x_i. The condition for diminishing marginal utility depends on parameters and on the partner's investment:

$$2\alpha^2\beta\phi(1 - x_j)\theta_i - 1 \leq 0$$

Therefore, diminishing returns holds globally for Firm i, i.e. for any precautions undertaken by its partner ($\chi \leq x_j \leq 1$), under the following maintained assumption:

$$\alpha^2\beta\phi(1 - \chi)\theta_i \leq \frac{1}{2}$$

Global diminishing returns is inherent in the exogenous contagion model ($\phi = 0$), but otherwise requires an upper bound on the loss.[6] The maintained required assumption simplifies the exposition that follows.[7]

IV. Best responses

Firm i's best response to x_j maximizes $u(\theta_i, x_i, x_j)$. The first-order condition for an interior solution ($\chi < x_i < 1$) is linear in x_i:

$$\alpha\{1-\beta[\alpha-\phi(1-\alpha)](1-x_j)\}\theta_i + 2\alpha^2\beta\phi(1-x_j)\theta_i x_i - x_i + \chi = 0$$

The necessary second-order condition for an interior solution follows from diminishing returns. Whether the solution is interior or at a corner depends on the usual complementary slackness conditions for the constraint $\chi \leq x_i \leq 1$. Zero marginal cost of security at the baseline level, however, implies that the best response is always positive ($x_i > \chi$):

$$b(x_j;\theta_i) \equiv \min\left\{\frac{\chi + \alpha\theta_i\{1-\beta(1-x_j)[\alpha-\phi(1-\alpha)]\}}{1-2\alpha^2\beta\phi\theta_i(1-x_j)}, 1\right\}$$

In the special case of exogenous contagion ($\phi = 0$), the best response curve is linear and upward sloping at an interior solution:

$$x_i = \min\left\{\chi + \alpha\theta_i\{1-\alpha\beta(1-x_j)\}, 1\right\}$$

More generally, the best response curve may be upward or downward sloping. Investments are strategic complements if the partners have mutually reinforcing

---

[6] The bound is necessary because total security is multiplicative in safety from the direct and indirect attacks. Consequently, the Firm i's expect benefit from precautions is a convex increasing function of x_i. In contrast, marginal benefit is constant if security were additive in safety from direct and indirect attacks, as in Acemoglu, Malekian, and Ozdaglar (2013).

[7] If diminishing returns fails, then the firm fully invests in security ($x_i = 1$). For some ranges in parameters this creates a discontinuity in the best response function, possibly compromising the existence of pure strategy equilibrium.

incentives to invest in precautions, i.e. if best response curves are upward sloping. Strategic substitutability is the opposite case of downward sloping best response curves.

Proposition 1: Investments are strategic substitutes if losses are sufficiently high (consistent with diminishing returns), except in the exogenous contagion model. Investments are strategic complements in the exogenous contagion model. Even in the balanced security model, investments are strategic complements if losses from breach are not too large, and if the probability of a direct attack is sufficiently large and baseline security is sufficiently low.

Proof: The slope of the best reply curve at an interior solution is

$$\frac{\alpha\beta\theta_i\{\alpha[1+\phi(1-2\chi)]-2\alpha^2\theta_i\phi-\phi\}}{[1-2\alpha^2\beta\theta_i\phi(1-x_j)]^2}$$

i.e. has the same sign as $\alpha-\phi[1-\alpha(1-2\chi)+2\alpha^2\theta_i]$, given diminishing returns. *(a) The sign is negative if* $\phi>0$ *and* θ_i *is sufficiently large. Recall that diminishing places an upper bound on* θ_i. *Consequently, there is a limited range of* θ_i *consistent with both diminishing returns and strategic substitutability:*

$$\frac{\alpha-\phi[1-\alpha(1-2\chi)]}{2\alpha^2\phi}\le\theta_i\le\frac{1}{2\alpha^2\beta\phi(1-\chi)}$$

This range of allowable θ_i *is non-degenerate for all* $0<\phi\le 1$. *(b) The sign is positive if* $1-\alpha(1-2\chi)+2\alpha^2\theta_i\le 0$, *or if* $1-\alpha(1-2\chi)+2\alpha^2\theta_i>0$ *and*

$$\phi\le\frac{\alpha}{1-\alpha(1-2\chi)+2\alpha^2\theta_i}$$

Furthermore, the inequality holds at $\phi=1$ *in the limit as* $\theta_i\to 0$ *if* $\alpha(1-\chi)>\frac{1}{2}$.
Q.E.D.

Strategic complementary thus depends on how difficult is to defend against indirect attacks. If security precautions serve mainly to defend against direct attacks, or if exposure to indirect attacks is low, then best response curves are upward sloping, and improved incentives for security are mutually reinforcing.

V. Equilibrium Security

In (Nash) equilibrium the partners correctly anticipate each other's investment. Equilibrium precautions occur at the intersection of best response curves, and the equilibrium is unique if the best response curves cross once. Uniqueness is readily established in the exogenous contagion case because the intercepts of the best response curves exceed χ, and the slopes at interior best responses are positive. The following uniqueness result, however, is more general.

Proposition 2: Equilibrium is unique.

Proof: First, as noted above, it is straightforward that $b(\chi;\theta_i) > \chi$. *Second, at an interior best response, the slope of the best response curve has the same sign as* $\alpha - \phi(1-\alpha+2\alpha^2\theta_i+2\alpha\chi)$. *Furthermore, given the second order condition (implied by diminishing returns), the best response curve is concave where it is strictly upward sloping and convex where it is strictly downward sloping. These properties imply that the best response curves intersect once. Q.E.D.*

The equilibrium best response could be at a corner, achieving complete security from direct attacks ($x_i = 1$), or interior ($0 < x_i < 1$). An interior equilibrium solves a fixed point to determine x_i:

$$x_i = b(b(x_i,\theta_j),\theta_i).$$

Since this equation can be transformed to a quadratic equation, an interior equilibrium can be solved in closed form.

Is the equilibrium security of a firm increasing in the magnitude of its loss from breach? The answer is not obvious if investments are strategic substitutes for the partner. Let $x^*(\theta_i,\theta_j)$ denote the equilibrium investment of Firm i, i.e. the fixed point of the above

equation determining x_i. Then the equilibrium security of Firm i is $s^*(\theta_i,\theta_j) = s(x^*(\theta_i,\theta_j), x^*(\theta_j,\theta_i))$. An increase in θ_i will increase $x^*(\theta_i,\theta_j)$ but might decrease $x^*(\theta_j,\theta_i)$ at the same time. Thus it seems theoretically possible that security of Firm i might go down if the strategic response of the partner is sufficiently strong. Numerical analysis, however, confirms that the direct effect on its own precautions dominates, and $s^*(\theta_i,\theta_j)$ is increasing in θ_i over parameter ranges satisfying diminishing returns.[8]

VI. Optimal security

A social planner maximizes the sum of utilities:

$$w(\theta_1,\theta_2,x_1,x_2) = u(\theta_1,x_1,x_2) + u(\theta_2,x_2,x_1)$$

The first-order condition with respect to x_i at an interior solution is the as same first-order equilibrium condition for Firm i plus an additional term that accounts for the externality:

$$\alpha\theta_i\{1-\beta[\alpha-\phi(1-\alpha)](1-x_j)\} + 2\alpha^2\beta\phi(1-x_j)\theta_i x_i + [1-\alpha(1-x_j)]\alpha\beta(1-\phi x_j)\theta_j - x_i + \chi = 0$$

In other words, the interior first-order conditions for the social planner are upward shifts of the best response curves:

$$\hat{b}(x_j;\theta_i,\theta_j) = \min\{\frac{\chi + \alpha\theta_i\{1-\alpha\beta(1-x_j)[1+\phi(1-\alpha)]\} + [1-\alpha(1-x_j)]\alpha\beta(1-\phi x_j)\theta_j}{1-2\alpha^2\beta\phi\theta_i(1-x_j)}, 1\}$$

This function defines a social best response curve.

---

[8] The parameters α, β, ϕ, and χ were each drawn uniformly from $[0,1]$, and, given these draws, θ_i and θ_j were each drawn uniformly from the intervals satisfying diminishing returns. For 1,000,000 such simulations, in no case did a 0.001 increased in θ_i increase $s^*(\theta_i,\theta_j)$.

Maintaining the assumption that the social welfare function is quasi-concave over the relevant range,[9] individual investments may be excessive or deficient compared to the optimum depending on the slopes of the best response curves and the asymmetry of losses.

Proposition 3: *If investments of both firms are strategic complements, then equilibrium investments providing less than complete security are deficient relative to the social optimum. If instead investments are strategic substitutes for Firm i, and expected losses are sufficiently asymmetric, then Firm i's equilibrium investment may be above the socially optimal level.*

Proof: *Figure 1 illustrates the case in which security investments are strategic complements for both firms, i.e. both best response curves are upward sloping. Since social best response curves lie above the (selfish) best response curves, it is obvious that the intersection of the social best response curves must occur at higher security investments than interior equilibrium investments. Figure 2 illustrates a mixed case in which investments are strategic complements for Firm 2 and strategic substitutes for Firm 1. For these particular parameter values, the intersection of the social best response curves is slightly to the left of the equilibrium point, i.e. in equilibrium Firm 1 overinvests relative to the first-best level. Q.E.D.*

[9] Numerical analysis confirms that the appropriate quasi-concavity conditions hold for most of the economically relevant parameter space if expected losses are not too large. Out of 1,000,000 points sampled uniformly from the parameter and strategy spaces, restricting θ_i to satisfy the diminishing returns conditions, there were 18.8% violations of the standard bordered Hessian condition for quasi-concavity. Typical violations are "edge cases", occurring if α, β, or ϕ is close to 0, or χ is close to 1, and θ is close to its upper bound.

Figure 1

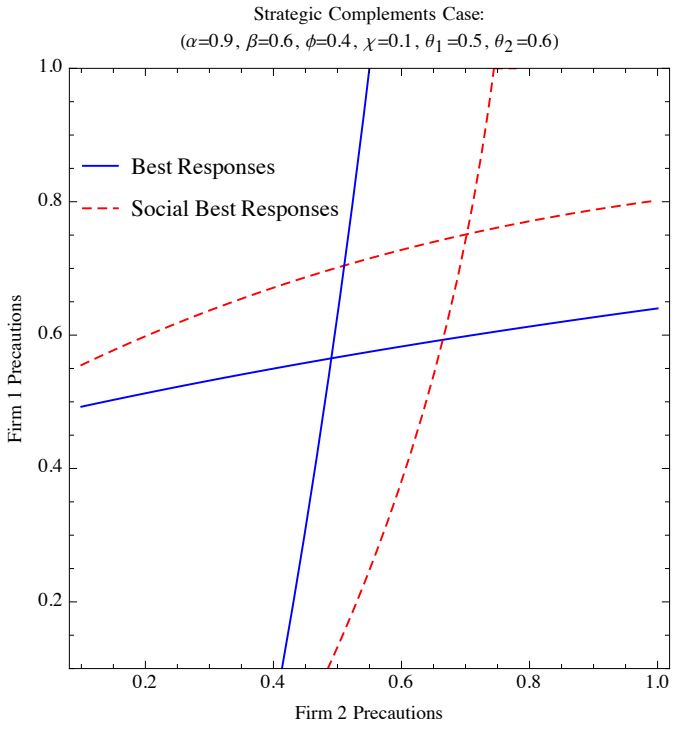

Figure 2

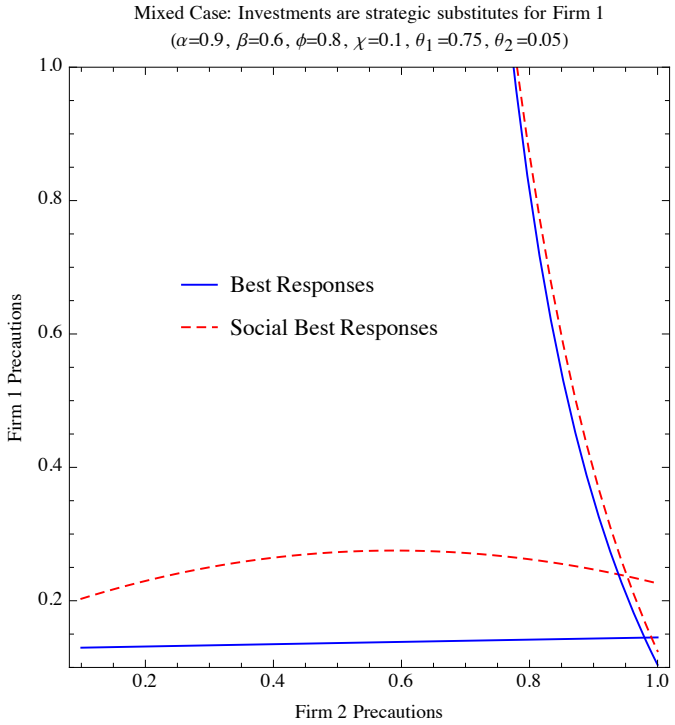

As discussed in the literature review, Acemoglu, Malekian, and Ozdaglar (2013) demonstrated a related excessive investment result due to asymmetric positioning in a balanced security network model. Here, deficient investment result is due instead to asymmetric losses. It important to note, however, that in any case investment is less than the social best response, i.e. each equilibrium investment is a socially deficient best response because of the positive externality.

Is optimal security increasing in the magnitude of losses? Let $\hat{x}(\theta_i,\theta_j)$ denote the optimal investment of Firm i. Then the optimal security of Firm i is $\hat{s}(\theta_i,\theta_j) = s(\hat{x}(\theta_i,\theta_j),\hat{x}(\theta_j,\theta_i))$. In contrast to the monotonicity of equilibrium security, numerical analysis reveals that $\hat{s}(\theta_i,\theta_j)$ is not necessarily increasing in θ_i. Most violations occur, however, at boundaries of the parameter space or for large θ_i.[10]

VII. Penalties and Bonuses

Requiring Firm i to pay fine equal to θ_j if its partner is breached by contagion solves the externality problem. A difficulty with this remedy is that it requires verification of the reason for a breach, which may not be contractible. A no-fault penalty however also achieves the first-best.

Proposition 4: Equilibrium coincides with the social optimum if Firm i is penalized by θ_j whenever Firm j is breached.

---

[10] Parameters α, β, ϕ, and χ were each drawn uniformly from the interval $[0.05,0.95]$, and θ_i and θ_j were each drawn uniformly from $[0,1]$. For 1,000,000 such simulations there were only 38 cases in which a 0.001 increase in θ_i increased $\hat{s}(\theta_i,\theta_j)$. These violations appear to be due to numerical imprecision.

Proof: If Firm i pays p whenever Firm j is breached, then Firm i's best response to x_j *is* $\hat{b}(x_j;\theta_i,p)$*, and therefore coincides with the social best response curve if* $p=\theta_j$*. Q.E.D.*

A virtue of an optimal no-fault fine is that it decentralizes investment decisions. However, implementation requires a third-party to collect the fines. Under a competitive contract, the third-party would pay to the partners upfront the equilibrium expected value of the fines. This solution is problematic if firms have limited liability. A limited liability constraint would put a bound on the net fine.

Limited liability constraints can be relaxed by implementing the first-best with a bonus scheme managed by a third-party with deep pockets. According to this scheme, Firm i would receive a bonus equal to θ_j if its partner does not suffer a breach, and make an upfront payment compensating the third-party for the expected value of the bonus. Limited liability constraints for the partners are relaxed because the third-party bears the risk of a high bonus payment.

Interpreting a foregone bonus as a penalty immediately establishes that an optimal bonus scheme fully internalizes externalities.

Corollary: Equilibrium coincides with the social optimum if Firm i is paid a bonus equal to θ_j *whenever Firm j is not breached.*

VIII. Standards

A standard can viewed as a set of best practices upon which parties can contract. Economics sees at least two potentially important roles for standards. The first is to enforce a floor on performance, e.g. a minimum quality standard. The second is to coordinate investments in complementary assets that require compatibility. The second

role seems less prominent in the context of cybersecurity. Therefore, the analysis below focuses on the first role.

Enforcement of a standard is likely to be less effective than well-crafted penalties and bonuses at improving equilibrium social welfare for two reasons. First, compliance with the standard must be verifiable. Second, there are inefficiencies inherent in enforcement of a standard. By definition, a standard is not customized to a particular situation, and, consequently, is likely to be more costly than decentralized investments that achieve the same level of security. Enforcement of a verifiable standard is welfare improving only if such inefficiencies are sufficiently small.

Suppose a contractible standard achieves security level γ at a cost $c(\gamma)+\eta(\gamma)$. The cost distortion $\eta(\gamma)$ is the additional cost of the standard compared to the cost of equally effective precautions. Thus, $\eta(\gamma)>0$ measures the inefficiency of standard compared to customized investments. It makes sense to assume that $\eta(\gamma)$ is increasing in γ, i.e. more ambitious standards are more distortionary.

Assume $\theta_1 \geq \theta_2$ and suppose that the standard is "modest" in the following two senses. First, the standard is less than the first-best investment for Firm 2. Second, if Firm 2 were to adopt the standard, then Firm 1 would have a private incentive for precautions that achieve a security level as least as good as the standard: $\gamma \leq b(\gamma,\theta_1)$. Under what conditions is it socially desirable to require only Firm 2 to adopt the standard technology? Under what conditions is it desirable to require both to adopt the standard?

<u>Proposition 5</u>: *Let x_2 denote Firm 2's equilibrium investment in the absence of a standard. Requiring Firm 2 to adopt a modest standard improves social welfare if* $\gamma > x_2$ and $\eta(\gamma)$ *is sufficiently small. Requiring Firm 1 also to adopt the standard reduces social welfare.*

Proof: *If Firm 2 adopts the standard, then Firm 1's best response is* $b(\gamma,\theta_1)$ *and social welfare is* $U(b(\gamma,\theta_1),\gamma,\theta_1)+U(\gamma,b(\gamma,\theta_1),\theta_2)-\eta$. *As* $\eta\to 0$, *the resulting welfare is above equilibrium welfare and below optimal welfare by definition of a modest standard. Requiring Firm 1 also to adopt the standard does not change incentives but incurs an additional cost* η. *Q.E.D.*

The optimal modest standard for the planner to impose on Firm 2 maximizes $W(b(\gamma,\theta_1),\gamma,\theta_1,\theta_2)-\eta(\gamma)$ subject to $\gamma\geq x_2$. The result will be a corner if $\eta(x_2)$ is sufficiently large and $\eta(\gamma)$ increases sufficiently quickly. Therefore, imposing a standard even selectively need not improve social welfare. The problem is that a poorly adapted standard imposes additional costs.

Standards for cybersecurity are fragmented and evolutionary. The exception is standards governing the handling of confidential personal information. The major credit card associations established the Payment Card Industry Data Security Standard (PCI DSS) to improve handling of confidential financial information. Unfortunately, compliance agreements did not prevent the Target and similar breaches. The verdict is still out on whether the problem is inadequate enforcement or insufficiency of the standard itself. Both Target and the company who certified Target's compliance were sued. At the same time, some commentators suggest that the encryption requirements of the standard may be insufficient.[11]

Another critical area in which the potential importance of cybersecurity standards has been recognized is the protection of critical infrastructure. The National Institute of Standards and Technology in 2014 released the Framework for Improving Critical Infrastructure Security. The document is a guide for systematic risk management by individual organizations. It itemizes published standards and best practices compiled by

---

[11] See (http://www.computerworld.com/s/article/9245709/_After_Target_Neiman_Marcus_breaches_does_PCI_compliance_mean_anything_?taxonomyId=203&pageNumber=2).

various standard-setting organizations. A potential virtue of this compilation of voluntary standards is that if provides a basis for contracting.

Markets for cybersecurity insurance are in a nascent state. While insurance is readily available for certain losses from confidential data breaches, insurance for losses from business disruption, reputational damage, and intellectual property theft generally is unavailable. Conditional insurance contracts are viewed as vehicle for promoting the adoption of standards that might improve cybersecurity overall (U.S. Department of Homeland Security, 2012).

Bonus contracts similarly could encourage the adoption of standards. The information requirements for bonus contracts are similar to those for insurance contracts. Furthermore, bonus contracts have the virtue of improving the underinvestment problem, whereas the moral hazard of insurance contracts potentially worsens the problem.

IX. Conclusions

The economic externality arising from the possibility of indirect attacks suggests that purely decentralized decision-making typically leads to deficient investments in security. Consequently, the efficient provision of security requires a centralized mechanism to correct externalities by penalizing for losses from security breaches. A bonus contract, by which a third-party rewards a firm if its partner is not breached, is a potentially attractive mechanism for improving interdependent security that can relax limited liability constraints, assuming the third-party has deep pockets.

A direction for future research is to relax the assumption that losses from breach are common knowledge. Incomplete information about losses introduces an additional organizational problem for the partnership. How can partners coordinate costly precautions to better manage interdependent security? Investments in precautions might be strategic substitutes, in the sense that greater investment by one party diminishes the

incentive of the other to invest. As we have seen, this is likely to be the case if security investments have strong public good attributes, in which case one party may find an incentive to free-ride on the other's private provision of the public good. Alternatively, investments might be strategic complements, so that increased investment by one party improves the other's incentive to invest. This is likely to be the case if it is difficult to improve protection against indirect attacks. In either case, coordination is complicated if security investments are difficult to observe and investment incentives are opaque. Similar principles apply for correcting the positive externality of security investments if the coordinating mechanism elicits truthful reports about expected losses.

Given the difficulties (e.g. limited liability) with using penalties and bonuses to correct incentives, and the additional complications introduced by incomplete information (e.g. about expected losses), an obvious remedy for the interdependent security problem is common ownership. In particular, cybersecurity could be an additional important force for a renewed reliance on vertical integration of manufacturing by original equipment manufactures (OEMs) who need to share proprietary technical information with specialized suppliers. Vertical integration has the benefit of directly internalizing the positive externality of security precautions by unifying the decision-making authority over security investments. Furthermore, the narrowing cost differences between offshore and onshore procurement, as well as the desire for more flexible manufacturing, is leading OEMs to consider the re-shoring of procurement.[12] In this context, improved cybersecurity plausibly could tip the balance in favor of greater vertical integration of the supply chain.[13]

---

[12] See, for example, "Why U.S. Manufacturing is Poised for a Comeback (Maybe)," *Wall Street Journal*, June 1, 2014.

[13] Loertscher and Riordan (2014) argue that the markup avoidance benefit of vertical integration is counterweighed by reduced incentives of independent suppliers to invest in vertical integration. If these primary forces are roughly in balance, then other factors, such as improved cybersecurity, could be decisive.

References

Daron Acemoglu, Azarakhsh Malekian, and Asu Ozdaglar, "Network Security and Contagion," working paper, 2013.

William J. Baumol, John C. Panzar, and Robert D. Willig (1982), *Contestable Markets and the Theory of Industry Structure*.

Jeremy I. Bulow, John D. Geanakoplos, and Paul D. Klemperer (1985), "Multimarket Oligopoly: Strategic Substitutes and Complements," *Journal of Political Economy*, 93, 488-511.

Lawrence A. Gordon and Martin P. Loeb, "The Economics of Information Security Investment," *ACM Transactions on Information and Security Systems*, Vol. 5 No. 4, November 2002, 438-457.

Geoffrey Heal and Howard Kunruether (2007), "Modeling Interdependent Risks," *Risk Analysis*, 27 (3), 621-634.

Howard Kunruether and Geoffrey Heal (2003), "Interdependent Security," *Journal of Risk and Uncertainty,* 26 (2/3), 231-249.

Simon Loertscher and Michael Riordan (2014), "Outsourcing, Vertical Integration, and Cost Reduction," preliminary working paper, 2014.

Mandiant, "MTrends: Attack the Security Gap, 2013 Threat Report," www.mandiant.com .

Symantec Corporation, "Internet Security Threat Report", Volume 18, 2013.

Hal R. Varian (2004), "System Reliability and Free Riding," in L.J. Camp and S. Lewis (eds.), *Economics of Information Security*, Kluwer Academic Publishers, 1-15.

U.S. Department of Homeland Security (2012), "Cybersecurity Insurance Workshop Readout Report."

U.S. Senate Committee on Commerce, Science, and Transportation (2014), "A 'Kill Chain' Analysis of the 2013 Target Data Breach," Majority Staff Report for Chairman Rockefeller, March 26.