
SUBSIDIES IN INTERDEPENDENT SECURITY WITH HETEROGENEOUS DISCOUNT RATES

Jun Zhuang and Vicki M. Bier

Department of Industrial and Systems Engineering,
University of Wisconsin–Madison, Madison, Wisconsin, USA

Ashish Gupta

Citigroup, Citi–Sears, Hoffman Estates, Illinois, USA

Optimizing resource allocation in interdependent security problems is a serious challenge for homeland security. In this article, we present the equilibrium strategies for multiple interdependent defenders in a model where threats occur over time. We show that the existence of myopic agents can make it undesirable for non-myopic agents to invest in security when it would otherwise be in their interests to do so. The phenomena of tipping and cascading are discussed, and we explore how to target subsidies for security investment in order to achieve the best results from tipping. The above findings are illustrated in numerical examples, and their policy implications are discussed.

INTRODUCTION

After September 11, 2001, homeland security has received a great deal of attention in the United States. Since security-related resources are always limited, maximizing security subject to limited resources is a key challenge. Many security problems (including aviation security, computer security, and supply-chain security) involve interdependence among potential defenders, meaning that one agent's strategy can affect the security environment for other agents. For example, poor security on the part of one airline, computer user, or supply-chain partner can increase the risk to other

Address correspondence to Jun Zhuang, Department of Industrial and Systems Engineering, University of Wisconsin–Madison, 1550 Engineering Drive, Room 3155, Madison, WI 53706. E-mail: jzhuang@wisc.edu

agents. Game theory has already been applied to such interdependent security problems. In particular, Heal and Kunreuther (2002) and Kunreuther and Heal (2003) discuss an interdependent security model in which agents are subjected both to direct attacks (which can be prevented by investment in security) and to indirect attacks (resulting from direct attacks on other agents, which we assume can be prevented only by the actions of those other agents).

However, the model of Kunreuther and Heal is static, in the sense that all attacks are assumed to occur at a single instant in time. Such a model does not adequately address real-world situations in which attacks happen over time. By contrast, in a time-dependent model, protection against a subset of attacks (such as direct attacks) would postpone the expected time until an attack but might not decrease the probability of succumbing to an attack eventually. This suggests that agents' discount rates may play a key role in determining equilibrium investment strategies. Thus, a revised analysis is needed to explain the effects of time preferences on the part of the agents in the model. In particular, this allows us to investigate the effects of heterogeneity in discount rates.

Kunreuther and Heal (2003) have noted that methods of coordination (such as regulation) are sometimes needed to help ensure that the optimum is achieved. However, it may be virtually impossible to enforce investment in security in some situations (especially when the number of agents is extremely large, as in computer security). In such cases, subsidizing or incentivizing non-investing agents in some fashion to encourage them to invest may be a reasonable alternative. In particular, "tipping" (Schelling, 1978; Gladwell, 2002) has been suggested as a cost-effective way to encourage security investment since, if successful, a subsidy or other incentive to encourage a relatively small number of agents to invest can induce other agents to also invest in security, which may be beneficial when investing is the social optimum.

One way of incentivizing investment is for security to be "bundled" with other goods and services (e.g., provided by one's Internet service provider). Another way in which non-investing agents might be given an incentive to invest in security would be if such investment were a requirement for lucrative federal contracts or for preferred-supplier relationships with industrial clients. (By analogy, federally funded academic researchers are often encouraged to engage in socially desirable practices such as educational outreach to underrepresented minorities—either as a requirement of funding or to increase the ratings of their proposals.) The analysis in this article shows that providing such incentives to a limited number of firms (e.g., government contractors) can make security investment sufficiently widespread that it becomes the norm even for firms that are not subject to such incentives.

In this article, incentives for security investment are modeled as outright subsidies. However, we believe that the results also apply to other types of incentives besides direct subsidies (e.g., bundling of security with other services or making security investment a requirement to be competitive for certain contracts).

We begin by formulating a general interdependent security model for an arbitrary number of agents, with attacks occurring over time. We then solve this model and identify dominant and equilibrium solutions—first for two agents, and then for arbitrary numbers of agents, focusing on the general case of agents with heterogeneous time preferences. We discuss the phenomena of tipping and cascading and determine the minimal number of agents that must receive subsidized security in order to cause tipping. We also discuss which agents should be targeted in order to maximize the beneficial effects of tipping and cascading. Finally, we discuss the policy implications of our work.

ASSUMPTIONS

Our basic model allows both direct attacks on an agent and indirect attacks (e.g., contamination from another agent in the system), as illustrated in Figure 1. We assume that the time, t , of a direct attack on any agent follows an exponential distribution, and that direct attacks can lead to indirect attacks with some probability.

Of course, in the real world, terrorist attacks do not occur randomly, but as the result of careful planning by attackers. However, considering attacks to be random with a constant rate of occurrence may not be unrealistic for some types of serious security threats, such as computer viruses. Therefore, for simplicity, our model views security solely as a game between defenders deciding whether to invest, and treats attacker behavior as exogenous. While this is obviously a somewhat limiting assumption,




Figure 1. Model structure.

we believe that our model will be valuable in generating insight into the relationships among defender choices and as a building block in developing more complex models that include endogenous attacker behavior as well as interactions among defenders. Models that treat security as a game between an intelligent attacker and a single defender include Bier et al. (in press), Konrad (2004), Woo (2002), Major (2002), Sandler and Arce (2003), Sandler and Lapan (1988), Lapan and Sandler (1993), and Zhuang and Bier (2006).

Like Kunreuther and Heal (2003), we assume that the loss is the same for both direct and indirect attacks, and that any attack is catastrophic (so that subsequent attacks affecting the same agent can be neglected). To justify the first assumption, we note, for example, that the loss from a computer virus is likely to be the same regardless of whether it is received directly from the virus developer or inadvertently from another infected computer. Similarly, the loss from exposure to an infectious bio-terrorism agent (such as smallpox or foot-and-mouth disease) may be the same regardless of whether the exposure is the result of an intentional attack or accidental contamination. To justify the assumption that even a single attack is catastrophic, we propose as examples situations in which a single successful attack might result in bankruptcy, death, loss of reputation, or theft of a valuable trade secret.

NOTATION AND MODEL FORMULATION

We define the system parameters as follows:

- N : Number of agents in the system, where $N \geq 2$.
- h : Number of agents receiving subsidized security, where $0 \leq h \leq N$.
- M : Number of agents having security measures (either choosing to invest or receiving subsidized security from a third party), where $0 \leq M \leq N$.
- x : Number of agents making erroneous choices, where $0 \leq x \leq N - h$.
- λ_i : Rate of direct attacks on agent i , for $i = 1, \dots, N$.
- $\tilde{\lambda}_i$: Total rate of all attacks on agent i (including indirect attacks), for $i = 1, \dots, N$.
- q_{ij} : Probability that an attack on agent i infects agent j (where we define $q_{ii} = 1$), for $i, j = 1, \dots, N$.
- r_i : Discount rate of agent i , where $r_i \geq 0$, for $i = 1, \dots, N$.
- L_i : Loss suffered by agent i if it is attacked, either directly or indirectly, for $i = 1, \dots, N$.

- C_i : Cost of investing in security for agent i . This investment is assumed to eliminate the risk of direct attacks, but have no effect on the risk of infection by indirect attacks from other agents. We assume that $0 < C_i < L_i \forall i = 1, \dots, N$.
- s_i : Investment strategy for agent i , for $i = 1, \dots, N$, where $s_i = 1$ if agent i invests in security, and $s_i = 0$ otherwise. We also denote S and N as $s_i = 1$ and $s_i = 0$, respectively.
- $s_{-i} \equiv \{s_j, j \neq i\}$: Set of strategies of all agents other than agent i .
- $P_i(s_i, s_{-i})$: Total expected cost borne by agent i , for $i = 1, \dots, N$, when it chooses strategy s_i (including both the cost of investment, if any, and the expected loss due to attacks), given the strategies of the other agents.

The expected loss experienced by agent i due to attacks is given by $L_i \int_0^\infty f_i(t) \exp(-r_i t) dt$, where $f_i(t) = \tilde{\lambda}_i \exp(-\tilde{\lambda}_i t)$ is the probability density function for the time of the first attack on agent i , and $\tilde{\lambda}_i = (1 - s_i)\lambda_i + \sum_{j \neq i} (1 - s_j)q_{ji}\lambda_j$ is the total rate of attacks against agent i . Hence, the net present value of the expected loss due to attacks experienced by agent i is

$$\begin{aligned}
 E(Loss) &= L_i \int_0^\infty \tilde{\lambda}_i \exp(-\tilde{\lambda}_i t - r_i t) dt \\
 &= \begin{cases} 0 & \text{if } \tilde{\lambda}_i = 0 \\ L_i / (1 + r_i / \tilde{\lambda}_i) & \text{if } \tilde{\lambda}_i > 0 \end{cases} \quad (1)
 \end{aligned}$$

and the total expected cost to agent i is given by

$$P_i(s_i, s_{-i}) = \begin{cases} s_i C_i & \text{if } \tilde{\lambda}_i = 0 \\ s_i C_i + L_i / (1 + r_i / \tilde{\lambda}_i) & \text{if } \tilde{\lambda}_i > 0 \end{cases} \quad (2)$$

DEFINITION 1. A pure-strategy Nash equilibrium (abbreviated as equilibrium in our paper) is a set of strategies $\{s_i, i = 1, \dots, N\}$ such that no one agent would be better off by switching strategies unless at least one other agent also switched. Thus, at equilibrium, the following system of inequalities must be satisfied:

$$P_i(s_i, s_{-i}) \leq P_i(1 - s_i, s_{-i}) \forall i = 1, \dots, N \quad (3)$$

DEFINITION 2. Strategy s_i is a (weakly) dominant strategy for agent i if and only if $P_i(s_i, s_{-i}) \leq P_i(1 - s_i, s_{-i})$ for all s_{-i} . Strategy s_i is a strictly dominant strategy for agent i if and only if $P_i(s_i, s_{-i}) < P_i(1 - s_i, s_{-i})$ for all s_{-i} .

For simplicity, in the remainder of this article, we consider the case where only the discount rates r_i of the agents differ. In other words, we let $\lambda_i = \lambda > 0$, $C_i = C > 0$, $L_i = L > 0$, and $q_{ij} = q > 0$ for $i, j = 1, \dots, N$, $i \neq j$, since the effects of these parameters have already been extensively investigated by Kunreuther and Heal (2003). Thus, Equation (2) becomes:

$$P_i(s_i, s_{-i}) = \begin{cases} s_i C & \text{if } \tilde{\lambda} = 0 \\ s_i C + L/(1 + r_i/\tilde{\lambda}_i) & \text{if } \tilde{\lambda} > 0 \end{cases} \quad (4)$$

where $\tilde{\lambda} = (1 - s_i)\lambda + \sum_{j \neq i} (1 - s_j)(q\lambda)$.

EQUILIBRIUM SOLUTIONS FOR TWO AGENTS

For simplicity, we begin our analysis with the case of $N = 2$. Table 1 calculates the total expected costs borne by agent $i = 1, 2$ for all four possible cases (with each agent either investing or not), using Equation (4). Investment in security will be a dominant strategy for agent i ($i = 1, 2$) when

$$C \leq \frac{L\lambda}{\lambda + r_i} \quad (5)$$

and

$$C \leq \frac{L\lambda r_i}{(\lambda + \lambda q + r_i)(\lambda q + r_i)} \quad (6)$$

Note that the right-hand side of (5) will always be greater than the right-hand side of (6), since the inequality

$$\frac{L\lambda}{\lambda + r_i} > \frac{L\lambda r_i}{(\lambda + \lambda q + r_i)(\lambda q + r_i)}$$

can be simplified to yield $\lambda^2 q + \lambda^2 q^2 + 2\lambda q r_i > 0$, which always holds. Hence, inequality (6) is the sufficient condition for investment to be the dominant strategy for agent i ; when this inequality holds, agent i will prefer to invest in security regardless of the investment decision made by the

Table 1. Costs for agents for $N = 2$

	Agent 2 invests $s_2 = 1$	Agent 2 does not invest $s_2 = 0$
Agent 1 invests $s_1 = 1$	$C; C$	$C + \frac{L\lambda q}{\lambda q + r_1}; \frac{L\lambda}{\lambda + r_2}$
Agent 1 invests $s_1 = 0$	$\frac{L\lambda}{\lambda + r_1}; C + \frac{L\lambda q}{\lambda q + r_2}$	$\frac{L(\lambda + \lambda q)}{\lambda + \lambda q + r_1}; \frac{L(\lambda + \lambda q)}{\lambda + \lambda q + r_2}$

other agent. However, even when inequality (6) is not satisfied, investing in security can still be an equilibrium strategy for agent i (although not a dominant strategy), depending on the investment decision of the other agent.

We now analyze the equilibrium strategies of the two agents for different values of r_1 and r_2 . There are four possible cases, as discussed below.

Case (A) Both Agents Invest in Security, (S, S)

(S, S) will be the unique equilibrium strategy when the cost of investing in security is low:

$$C \leq \min_{i=1,2} \left[\frac{L\lambda}{\lambda + r_i} \right];$$

and also

$$C \leq \max_{i=1,2} \left[\frac{L\lambda r_i}{(\lambda + \lambda q + r_i)(\lambda q + r_i)} \right]$$

Case (B) Neither Agent Invests in Security, (N, N)

(N, N) will be the unique equilibrium strategy when the cost of investing in security is high:

$$C \geq \max_{i=1,2} \left[\frac{L\lambda r_i}{(\lambda + \lambda q + r_i)(\lambda q + r_i)} \right];$$

and also

$$C \geq \min_{i=1,2} \left[\frac{L\lambda}{\lambda + r_i} \right]$$

Case (C) One Agent Invests but the Other Does Not, (N, S) or (S, N)

There will be a unique Nash equilibrium strategy, either (N, S) or (S, N), when the cost of investing in security is high relative to the expected net present value of the loss due to direct attack for one agent, $C \geq L\lambda/(\lambda + r_i)$, but low relative to the increased loss of the other agent, $C \leq L\lambda r_j/((\lambda + \lambda q + r_j)(\lambda q + r_j))$, where $i \neq j$. In this situation, the decisions of the two agents are independent of each other. Agent i (with the higher discount rate) will prefer not to invest in security, while agent j (with the lower discount rate) will prefer to invest regardless of the actions of the other agent.

Case (D) Both Agents Either Invest or Don't Invest in Security, (S, S) or (N, N)

This game can have multiple equilibrium strategies, (S, S) and (N, N), when the cost of investing in security is intermediate for both agents: $L\lambda r_i / ((\lambda + \lambda q + r_i)(\lambda q + r_i)) \leq C \leq L\lambda / (\lambda + r_i)$, nor $i = 1, 2$. In this case, neither agent has a dominant strategy; instead, each agent will prefer to adopt whichever strategy is chosen by the other agent. Thus, if one agent decides to invest, the other agent will also find it profitable to invest, and conversely. Of these two equilibrium strategies, we would prefer for the one that is socially optimal to be adopted. Investment in security will be the socially optimal decision when $C \leq L(\lambda + \lambda q) / 2(1 / (\lambda + \lambda q + r_1) + 1 / (\lambda + \lambda q + r_2))$.

Table 2 summarizes all possible equilibrium strategies, where for simplicity, we let $a_i = L\lambda r_i / ((\lambda + \lambda q + r_i)[\lambda q + r_i])$ and $b_i = L\lambda / (\lambda + r_i)$. To see the implications of the agents' discount rates on their equilibrium strategies, we also show the equilibrium strategies as a function of the agents' discount rates instead of the investment cost.

When agent i has a high discount rate, $r_i \geq \lambda(L/C - 1)$, future losses due to attacks will have a low present value, so agent i will not find it worthwhile to invest in security. Similarly, when agent i has a moderately small discount rate,

$$\frac{\lambda}{2} \left\{ \left(\frac{L}{C} - 1 \right) - 2q - \sqrt{\left(\frac{L}{C} - 1 \right)^2 - 4\frac{qL}{C}} \right\} \leq r_i \leq \frac{\lambda}{2} \left\{ \left(\frac{L}{C} - 1 \right) - 2q + \sqrt{\left(\frac{L}{C} - 1 \right)^2 - 4\frac{qL}{C}} \right\}$$

the losses due to future attacks will tend to loom relatively large, so agent i will find investing in security to be worthwhile. When the discount rate

Table 2. Equilibrium strategies as a function of the investment cost

	$C \leq a_2$	$a_2 \leq C \leq b_2$	$C \geq b_2$
	(S, S)	(S, S)	(S, N)
$C \leq a_1$	Dominant strategy	Nash equilibrium	Nash equilibrium
	(S, S)	(S, S) or (N, N)	(N, N)
$a_1 \leq C \leq b_1$	Nash equilibrium	Multiple Nash equilibria	Nash equilibrium
	(N, S)	(N, N)	(N, N)
$C \geq b_1$	Nash equilibrium	Nash equilibrium	Dominant strategy

of agent i is in the intermediate range,

$$\frac{\lambda}{2} \left\{ \left(\frac{L}{C} - 1 \right) - 2q + \sqrt{\left(\frac{L}{C} - 1 \right)^2 - 4 \frac{qL}{C}} \right\} \leq r_i \leq \lambda \left(\frac{L}{C} - 1 \right)$$

agent i will be ambivalent about whether to invest. If the other agent decides to invest (and hence eliminates the risk of indirect losses to agent i), investing will become more cost effective for agent i , since investing will now eliminate all risk, rather than only a subset of the total risk. However, if the other agent decides not to invest, investing will no longer be cost effective for agent i , since it can prevent only direct losses, not losses due to indirect attacks. Thus, agent i will prefer to adopt the same strategy as the other agent in this case.

Finally, when the discount rate of agent i is extremely small,

$$0 \leq r_i \leq \frac{\lambda}{2} \left\{ \left(\frac{L}{C} - 1 \right) - 2q - \sqrt{\left(\frac{L}{C} - 1 \right)^2 - 4 \frac{qL}{C}} \right\}$$

eliminating only a subset of the total risk is no longer worthwhile for agent i , since it only postpones the loss from an attack rather than eliminating it. Therefore, in this situation, it is again worthwhile for agent i to invest in security only if the other agent also invests.

For simplicity, we define the following:

$$\rho_a = \frac{\lambda}{2} \left\{ \left(\frac{L}{C} - 1 \right) - 2q - \sqrt{\left(\frac{L}{C} - 1 \right)^2 - 4 \frac{qL}{C}} \right\};$$

$$\rho_b = \frac{\lambda}{2} \left\{ \left(\frac{L}{C} - 1 \right) - 2q + \sqrt{\left(\frac{L}{C} - 1 \right)^2 - 4 \frac{qL}{C}} \right\};$$

and

$$\rho_c = \lambda \left(\frac{L}{C} - 1 \right),$$

where $\rho_a < \rho_b < \rho_c$. With this notation, Table 3 summarizes all possible equilibrium strategies as a function of the agents' discount rates r_i .

Table 3. Equilibrium strategies as a function of the discount rates for $N = 2$

r_1	r_2			
	Extremely small ($0 \leq r_2 \leq \rho_a$)	Moderately small ($\rho_a \leq r_2 \leq \rho_b$)	Intermediate ($\rho_b \leq r_2 \leq \rho_c$)	High ($r_2 \geq \rho_c$)
Extremely small ($0 \leq r_1 \leq \rho_a$)	(S, S) or (N, N) Multiple Nash equilibria	(S, S) Nash equilibrium	(S, S) or (N, N) Multiple Nash equilibria	(N, N) Nash equilibrium
Moderately small ($\rho_a \leq r_1 \leq \rho_b$)	(S, S) Nash equilibrium	(S, S) Dominant strategy	(S, S) Nash equilibrium	(S, N) Nash equilibrium
Intermediate ($\rho_b \leq r_1 \leq \rho_c$)	(S, S) or (N, N) Multiple Nash equilibria	(S, S) Nash equilibrium	(S, S) or (N, N) Multiple Nash equilibria	(N, N) Nash equilibrium
High ($r_1 \geq \rho_c$)	(N, N) Nash equilibrium	(N, S) Nash equilibrium	(N, N) Nash equilibrium	(N, N) Dominant strategy

Equilibrium Solutions for Arbitrary Numbers of Agents

To characterize the more general game for $N \geq 3$, suppose that in some equilibrium solution, exactly M agents receive subsidized security or choose to invest (for some $0 \leq M \leq N$). Then there are three possible cases:

- All agents have security measures in place ($M = N$);
- Some agents have security measures and some not ($1 \leq M \leq N - 1$);
or
- No agents have security measures in place ($M = 0$).

For convenience, we renumber the agents so that the first M agents have security measures, and the remaining $N - M$ agents do not (i.e., $s_i = 1 \forall i = 1, \dots, M$, and $s_i = 0 \forall i = M + 1, \dots, N$). We begin by considering a model that does not include the possibility of subsidized security (i.e., with $h = 0$). Table 4 gives the costs for the agents in all three of the above cases for that model. It also specifies the cost to any agent i that deviates from the equilibrium action s_i , conditional on the other agents' strategies; i.e., $P_i(1 - s_i, s_{-i})$.

Solving the system of inequalities (3) for $i = 1, \dots, N$ using the costs in Table 4 yields conditions for the discount rates at which agents will be willing to invest in security at equilibrium. These conditions are summarized in Table 5, using the following notation:

Table 4. Costs for agents when M agents having security

	For investing agents ($i = 1, \dots, M$), $s_i = 1$	For not-investing agents ($i = M+1, \dots, N$), $s_i = 0$
	$P_i(s_i, s_{-i})^*$	$P_i(1 - s_i, s_{-i})$
Case (1) $M = N$	C	N/A
Case (2) $1 \leq M \leq N - 1$	$C + L/(1 + r_i/\tilde{\lambda})$ $\tilde{\lambda} = \lambda + (N - M)q\lambda$	$L/(1 + r_i/\tilde{\lambda})$ $\tilde{\lambda} = \lambda + (N - M - 1)q\lambda$
Case (3) $M = 0$	N/A	$L/(1 + r_i/\tilde{\lambda})$ $\tilde{\lambda} = \lambda + (N - 1)q\lambda$

*The costs for subsidized agents are the same with $P_i(s_i, s_{-i})$ of investing agents subtracting the investment cost C .

Table 5. Conditions for an equilibrium in which M agents have security

	For investing agents ($i = 1, \dots, M$), $s_i = 1$	For not-investing agents ($i = M+1, \dots, N$), $s_i = 0$
Case (1): $M = N$	$R_1(0) \leq r_i \leq R_2(0)$	N/A
Case (2):* $1 + (N - \tilde{N})^+ \leq M \leq N - 1$	$R_1(N - M) \leq r_i \leq R_2(N - M)$	$r_i \leq R_1(N - M - 1)$ or $r_i \geq R_2(N - M - 1)$
Case (3): $M = 0$	N/A	$r_i \leq R_1(N - 1)$ or $r_i \geq R_2(N - 1)$

*No equilibrium is possible with $1 \leq M \leq (N - \tilde{N})^+$.

- $\tilde{N} \equiv \lfloor C(L/C - 1)^2 / (4Lq) \rfloor$ (where $\lfloor x \rfloor$ is the greatest integer less than or equal to x);
- $R_1(k) \equiv \lambda[L/C - 1 - 2kq - \sqrt{(L/C - 1)^2 - 4qkL/C}] / 2$ for $k = 0, \dots, \tilde{N}$; and
- $R_2(k) \equiv \lambda[L/C - 1 - 2kq + \sqrt{(L/C - 1)^2 - 4qkL/C}] / 2$ for $k = 0, \dots, \tilde{N}$.

Here, \tilde{N} is a bound on the number of agents there can be in a system for certain properties to hold, and $R_1(k)$ and $R_2(k)$ are the minimum and maximum discount rates for which an agent would invest given that exactly k others are not doing so, respectively. By taking derivatives, it is easy to show that $R_1(k)$ and $R_2(k)$ are increasing and decreasing in k , respectively. Also, note that $R_1(\tilde{N}) \leq R_2(\tilde{N})$, and $R_1(0) = 0$. Thus, we have

$$0 = R_1(0) < R_1(1) < \dots < R_1(\tilde{N}) \leq R_2(\tilde{N}) < \dots < R_2(1) < R_2(0)$$

as shown in Figure 2.

We define the following sets on the domain of discount rates:

- $S_l = (R_2(0), \infty)$ for $l = 0$
- $S_l = (R_1(l - 1), R_1(l)) \cup (R_2(l), R_2(l - 1))$ for $l = 1, \dots, \min(N - 1, \tilde{N})$, and
- $S_l = (R_1(l - 1), R_2(l - 1))$ for $l = \min(N, \tilde{N} + 1)$




Figure 2. Illustration of the ranges S_l .

Intuitively, S_l is the set of those discount rates for which an agent would want to invest if at most $l - 1$ others do not invest but would not want to invest if l or more others do not invest. This notation is represented graphically in Figure 2 for the case where $N \leq \tilde{N} + 1$. Note that for $N = 2$, we have $R_1(1) = \rho_a$, $R_2(1) = \rho_b$ and $R_2(0) = \rho_c$.

THEOREM 1. *Holding all else constant, as the number of agents N in the system increases, the range of discount rates S_N for which investing is a strictly dominant strategy becomes smaller. If $N \geq \tilde{N} + 2$, then there is no discount rate for which investing in security is a strictly dominant strategy. The range of discount rates S_0 for which not investing is a strictly dominant strategy does not depend on N .*

PROOF. According to Table 5, if $r_i \in S_N$ for some agent i , then not investing can never be a best response for that agent, no matter how many or how few other agents choose to invest. Therefore, investing must be a strictly dominant strategy for agent i . Similarly, if $r_i \in S_0$, then not investing must be a strictly dominant strategy. Since $R_1(k)$ and $R_2(k)$ are increasing and decreasing in k for $k = 1, \dots, \tilde{N}$, respectively, the set $S_N = [R_1(N - 1), R_2(N - 1)]$ becomes smaller as N increases. For $N \geq \tilde{N} + 2$, the set S_N is empty. Finally, the set $S_0 = (R_2(0), +\infty)$ does not depend on N , by the definition of $R_2(0)$.

REMARK. This theorem implies that for systems with sufficiently large numbers of agents, investing in security will not be a strictly dominant strategy for any agent. Agents can of course still choose to invest, but will do so in equilibrium only if other agents also have security in place. This finding helps to justify the need for coordinating mechanisms (such as subsidies).

EXAMPLE 1. We use the following parameter values to illustrate the results in Theorem 1: $C = 10$; $L = 1000$; $q = 0.01$; $\lambda = 0.01$; $N = 2000$; and therefore $\tilde{N} = 2450$. Figure 3 shows the regions of discount rates in which investing and not investing, respectively, are dominant strategies, as functions of the number of agents.

Tipping and Cascading

We now discuss the possibility of tipping, and its effect on the equilibrium solution. In general, starting with an equilibrium in which M agents invest (for $M \leq N - 1$), ensuring that some additional agents invest will tend to make investing more attractive for the remaining agents. In practice, there may be various ways of encouraging those agents to invest, such as mandating investment in security, or providing subsidized security. In this




Figure 3. Discount rates for which investing and not investing are dominant, as a function of N .

article, we focus on the latter case and assume in particular that security becomes free to some agents.

DEFINITION 3. Given a system in which exactly M agents have security measures in place at equilibrium:

- Let $Inv(M)$ be the set of possible discount rates for the investing agents (note that there may be fewer than M investing agents if subsidized agents are also considered. However, those agents actually investing would still need to have discount rates in $Inv(M)$ if a total of M agents have security in place, regardless of whether some of those M agents obtained their security through subsidies);
- Let $Non(M)$ be the set of possible discount rates for the (non-subsidized) agents that do not invest;
- Let $Cmpl(M) = [Inv(M) \cup Non(M)]^C$ be the set of discount rates that can not be held by any (non-subsidized) agent if exactly M agents have security measures at equilibrium; and




Figure 4. Ranges of discount rates possible for investing agents when M agents invest or receive subsidized security.

- Let $Inv_c(M, h)$ be the set of possible discount rates for those agents who would not invest at equilibrium with no subsidy, but would find investing attractive if the number of other agents having security measures in place increased from M to $M + h$.

The ranges $Inv(M)$, $Non(M)$ and $Cmpl(M)$ are illustrated in Figures 4 and 5 for the investing and non-investing agents, respectively. Note that the existence of an equilibrium strategy with M investing agents implies that no (non-subsidized) agent has a discount rate in the set $Cmpl(M)$.

Note also that the set $Inv_c(M, h)$ contains the set $Cmpl(M + h)$, since there can be an equilibrium strategy with at least $M + h$ agents having security measures only if no other non-subsidized agents have discount rates in the set $Cmpl(M + h)$. Thus, we will have

$$Inv_c(M, h) = [Inv(M + h) - Inv(M)] \cup Cmpl(M + h)$$

Here, “-” is the set operator representing the difference between the sets $Inv(M + h)$ and $Inv(M)$. If we let $\Theta(S)$ denote the number of agents with discount rates in the set S , then $\Theta[Inv_c(M, h)]$ is the number of agents




Figure 5. Ranges of discount rates possible for non-investing agents when M agents invest or receive subsidized security.

that could be induced to invest by tipping, where $\Theta[Inv_c(M,h)]$ is non-decreasing in h for any given value of M , with $\Theta[Inv_c(M,0)] = 0$.

Starting with an equilibrium in which exactly M agents invest, the number of agents h that must receive subsidized security in order to lead to tipping must satisfy $\Theta[Inv_c(M,h)] > 0$. Moreover, if $\Theta\{Inv_c[M+h, \Theta[Inv_c(M,h)]]\} > 0$, then additional agents will choose to invest, in which case the number of agents having security measures will increase from $M+h$ to $M+h+\Theta[Inv_c(M,h)]$, and so on. We call this phenomenon “cascading.”

No tipping or cascading will occur if $\Theta[Inv_c(M,h)] = 0$. The minimal number of agents that must receive free security in order to lead to tipping is given by $\min\{h : \Theta[Inv_c(M,h)] > 0\}$. The discount rates of the subsidized agents are irrelevant to determining whether tipping occurs. However, the discount rates of the non-subsidized agents do determine whether cascading occurs, and how far it progresses. Therefore, it makes sense if possible to target any subsidies at those agents who are least likely to begin investing without such subsidies (i.e., agents with discount rates in or near the region where not investing is the strictly dominant strategy). These arguments are illustrated below; see also Kunreuther and Heal (2003) and Dixit (2003).

EXAMPLE 2. Consider an N -agent system with $h = 0$, $N \leq \tilde{N} + 1$, and $r_i \in S_{i-1} \forall i = 1, \dots, N$. Perusal of Figure 4 indicates that for this example, there exists no equilibrium with a non-zero number of agents investing. Furthermore, Figure 5 shows that the solution with $M = 0$ agents investing, $\{\text{none invest}\}$, is an equilibrium for this example. Now suppose that agent 1 receives free security. Then agent N will also choose to invest, because $r_N \in S_{N-1} \subset Inv_c(0,1)$. In other words, since agent 1 is receiving subsidized security, not investing is no longer optimal for agent N (from Figure 5). Since agent N is better off investing, there will now be $M = 2$ agents having security measures. Therefore, agent $N - 1$ will also begin investing, because $r_{N-1} \in S_{N-2} \subset Inv_c(1,1)$. Similarly, agent $N - 2$ will begin investing once there are $M = 3$ agents having security measures, and so on. Thus, if agent 1 receives free security, $\{\text{all invest}\}$ becomes the unique equilibrium in this example. Note also that the discount rate of the subsidized agent determines how many agents will decide to invest as a result of cascading. In this example, it is straightforward to see that if agent i is the one that receives free security, then the system will end up in a unique equilibrium with $M = N + 1 - i$ agents adopting security measures.

CONCLUSIONS

In the real world, agents may have different discount rates for a variety of reasons. First, firms in different industries (with differing levels of risk)

or with differing conditions of their physical plant (and hence different opportunities for investing in modernization) will tend to apply differing minimum acceptable rates of return in their decisions (Grant et al., 1998; Fabrycky et al., 1998). This can create challenges for supply-chain partnerships that cross industry lines. For example, a computer manufacturer or software vendor (in a high-risk industry with a high minimum acceptable rate of return) may invest less in security than some of its clients (say, electric utilities, in an industry with a historically low rate of return) would prefer.

Secondly, firms facing impending bankruptcy may have high discount rates (Golbe, 1981) and be willing to take much greater risks (including security risks) than firms that are financially stable. Woo (2004) has pointed out that much the same logic may apply to financially or otherwise unstable nation states with regard to the risk of terrorism.

Finally, of course, agents may simply be myopic (Kunreuther et al., 1998; Thaler et al., 1997; Langer and Weber, 2005), in the sense of adopting higher effective discount rates and shorter time horizons than would be in their own (enlightened) self-interest. The fact that the existence of such myopic players can eliminate the incentives for more prudent agents to invest in security is analogous to the results of Smith et al. (1988), who found that a non-zero probability of there being an irrational agent in a market can make investment in speculative bubbles a rational choice for other agents.

Thus, extending the model of Heal and Kunreuther (2002) and Kunreuther and Heal (2003) to one in which attacks occur stochastically over time shows that heterogeneity in discount rates can complicate the task of achieving rational security strategies in an interdependent world. Recognition of this phenomenon will hopefully contribute to both an improved understanding of the security challenges that we face and an enhanced ability to identify promising solutions to those challenges.

This article formulates and solves an interdependent security model for an arbitrary number of defenders with attacks occurring over time, focusing on the case of agents with heterogeneous time preferences. Results show that when multiple equilibrium solutions exist, the social optimum is for all agents to invest, as long as that is an equilibrium solution.

The role of tipping and cascading in helping to achieve this social optimum is also discussed. In particular, we explore the minimal number of agents who would need to receive subsidized (free) security in order for tipping to occur and which agents should ideally receive such free security. (However, we recognize that the mechanisms by which subsidies could be implemented in the real world may not allow such careful targeting.) More generally, our results emphasize the potential importance of coordination mechanisms—not only subsidies and incentives, but also regulation, trade associations, social norms, contracts, insurance, etc. Moreover, our results

suggest that attempts to encourage investment should be targeted at those agents least likely to begin investing on their own—which in general will be those with discount rates that are either extremely low or (perhaps more realistically) extremely high.

In this article, we focus primarily on the effect of discount rates; in particular, we show that the mere existence of agents with extreme discount rates (e.g., due to myopia) can make it undesirable for other agents to invest. However, we expect that similar results would also hold for heterogeneity in other parameters (such as the cost of investing in security, or the loss resulting from a successful attack).

ACKNOWLEDGMENTS

This material is based upon work supported in part by the U.S. Army Research Laboratory and the U.S. Army Research Office under grant number DAAD19-01-1-0502, by the U.S. National Science Foundation under grant number DMI-0228204, and by the U.S. Department of Homeland Security through the Center for Risk and Economic Analysis of Terrorism Events (CREATE) under grant number N00014-05-0630. Any opinions, findings, and conclusions or recommendations expressed herein are those of the authors and do not necessarily reflect the views of the sponsors. We also thank Larry Samuelson of the University of Wisconsin–Madison for his helpful comments and suggested references.

REFERENCES

- Bier, V.M., Oliveros, S. and Samuelson, L. (in press) Choosing what to protect. *Journal of Public Economic Theory*.
- Dixit, A.K. (2003) Clubs with entrapment. *American Economic Review*, 93, 1824–1829.
- Fabrycky, W.J., Thuesen, G.J. and Verma, D. (1998) *Economic Decision Analysis*, 3rd ed. Upper Saddle River, NJ: Prentice-Hall.
- Gladwell, M. (2002) *The Tipping Point: How Little Things Can Make a Big Difference*. Boston, MA: Back Bay Books.
- Golbe, L.D. (1981) The effects of imminent bankruptcy on stockholder risk preferences and behavior. *The Bell Journal of Economics*, 12, 321–328.
- Grant, E.L., Ireson, W.G. and Leavenworth, R.S. (1990) *Principles of Engineering Economy*, 8th ed. Hoboken, NJ: Wiley.
- Heal, G. and Kunreuther, H. (2002) You can only die once: Public-private partnerships for managing the risks of extreme events. White paper for conference on “Risk Management Strategies in an Uncertain World,” Palisades, New York. Available at: <http://www2.gsb.columbia.edu/faculty/gheal/General%20Interest%20Papers/onlydieonce.pdf> (accessed January 2007)
- Konrad, K.A. (2004) The investment problem in terrorism. *Economica*, 71, 449–459.

- Kunreuther, H. and Heal, G. (2003) Interdependent security. *Journal of Risk and Uncertainty*, 26, 231–249.
- Kunreuther, H., Onculer, A. and Slovic, P. (1998) Time insensitivity for protective investments. *Journal of Risk and Uncertainty*, 16, 279–299.
- Langer, T. and Weber, M. (2005) Myopic prospect theory vs. myopic loss aversion: How general is the phenomenon? *Journal of Economic Behavior and Organization*, 56, 25–38.
- Lapan, H.E. and Sandler, T. (1993) Terrorism and signaling. *European Journal of Political Economy*, 9(3), 383–397.
- Major, J. (2002) Advanced techniques for modeling terrorism risk. *Journal of Risk Finance*, 4, 15–24.
- Sandler, T. and Arce M.D.G. (2003) Terrorism and game theory. *Simulation & Gaming*, 34, 319–337.
- Sandler, T. and Lapan, H.E. (1988) The calculus of dissent: An analysis of terrorists' choice of targets. *Synthese*, 76, 245–261.
- Schelling, T. (1978) *Micromotives and Macrobehavior*. New York: Norton.
- Smith, L.V., Suchanek, L.G. and Williams, W.A. (1988) Bubbles, crashes, and endogenous expectations in experimental spot asset markets. *Econometrica*, 56, 1119–1151.
- Thaler, H.R., Tversky, A., Kahneman, D. and Schwartz, A. (1997) The effect of myopia and loss aversion on risk taking: An experimental test. *The Quarterly Journal of Economics*, 112, 647–661.
- Woo, G. (2002) Quantitative terrorism risk assessment. *Journal of Risk Finance*, 4, 7–14.
- Woo, G. (2004) Personal communication.
- Zhuang, J. and Bier, V.M. (2006) Balancing terrorism and natural disasters—Defensive strategy with endogenous attack effort. *Operations Research*.

BIOGRAPHICAL SKETCHES

JUN ZHUANG is a Ph.D. student in the Department of Industrial and Systems Engineering at the University of Wisconsin–Madison. His current research interests include the application of operations research, decision analysis, risk analysis, and game theory to problems of homeland security and supply-chain management.

VICKI M. BIER is a professor in the Department of Industrial and Systems Engineering at the University of Wisconsin–Madison. Her current research interests include the application of game theory to identify optimal resource allocation strategies for protecting critical infrastructure from intentional attacks. Other interests include the use of accident “precursors” or near misses in probabilistic risk analysis, the use of expert opinion, and methods for effective risk communication, both to decision-makers and to the general public.

ASHISH GUPTA is a 2005 M.S. graduate of the Department of Industrial and Systems Engineering at the University of Wisconsin–Madison. He is now with Citigroup, Citi-Sears.