Interdependent Security

Howard Kunreuther
Wharton School
University of Pennsylvania
Philadelphia, PA 19104
E-mail: kunreuther@wharton.upenn.edu

Geoffrey Heal
Graduate School of Business
Columbia University
New York, NY 10027
E-mail: gmh1@columbia.edu

December 2002

Journal of Risk and Uncertainty (forthcoming)

We acknowledge support for this research from the U.S. Environmental Protection Agency under Cooperative Agreement C R 826583 with the University of Pennsylvania, the Wharton Risk Management and Decision Processes Center and the Columbia University Earth Institute. We are particularly grateful to Richard Zeckhauser for helpful discussion and comments on earlier drafts. Useful insights were also provided by Stan Baiman, Mark Broadie, David Croson, Rachel Croson, Ido Erev, Victor Goldberg, Jay Hamilton, Jack Hershey, Daniel Kahneman, Paul Kleindorfer, Erwann Michel-Kerjan, Felix Oberholzer-Gee, Yechiam Yemini and participants in the NBER Insurance Project Workshop and workshops at Columbia University, Princeton University, Stanford University and the University of Pennsylvania.

ABSTRACT

Do firms have adequate incentives to invest in protection against a risk whose magnitude depends in the actions of others? This paper characterizes the Nash equilibria for this type of interaction between agents, which we call the *interdependent security (IDS)* problem. When agents are identical, there are two Nash equilibria for a wide range of cost and risk parameters --- either everyone invests in protection or no one does. In some situations the incentive to invest in protection approaches zero as the number of unprotected agents increases. We develop an IDS model by first focusing on airline security and comparing the structure of this problem with other IDS examples such as computer security, fire protection, vaccinations, protection against bankruptcy, and theft protection. The paper also examines the roles of insurance, liability, fines and subsidies, third party inspections, regulations and coordinating mechanisms for internalizing the negative externalities characteristic of these problems. The concluding section suggests directions for future theoretical and empirical research.

Key Words: externalities; contagion; protection; terrorism; Nash equilibrium

JEL Classification: C7 H2 D62

Interdependent Security

Howard Kunreuther
Geoffrey Heal

1. Introduction

In today's world of terrorist threats, many individuals and organizations are considering whether to invest more in security precautions. Do individuals and firms have incentives to carry out socially appropriate levels of security investment? Or are there reasons to suppose that they will systematically underinvest in this area? Ayres and Levitt (1998) have demonstrated the social benefits of protection when individuals invest in unobservable precautionary measures. They focus on the LoJack car retrieval system that criminals cannot detect. This generates positive externalities that naturally lead to a sub-optimal level of private investment.

This paper also focuses on situations where the security levels of members of a group are interdependent, and when investing in protection produces positive externalities. However, in contrast to expenditures on crime protection, the incentives to invest in security may be perverse: the dependence of one agent's security on the behavior of others may partially or in some cases almost completely negate the payoffs it receives from its own investment in protective measures. We refer to these cross-effects between one agent's incentives and the behavior of the others as *contagion*.

In using this word we draw an analogy with the phenomena studied in the literature on financial contagion where the issue is that a perceived financial weaknesses in one institution can lead to weaknesses in others that were not initially vulnerable [Musumeci and Sinkey (1990), Poloncheck and Miller (1999) and Allen and Gale (2000)]. In such situations each institution's vulnerability depends not only on the way in which it manages its risks but also on the ways in which other unrelated entities manage their risks. This is a similar structure to that studied here.

We illustrate the general argument by reference to an airline that is determining whether to install a baggage checking system voluntarily. In making this decision it needs to balance the cost of installing and operating such a system with the reduction in the risk of an explosion from a piece of luggage not only from the passengers who check in with it, but from the bags of passengers who check in on other airlines and then transfer to it.

The incentive to invest in security is greatly diminished if other airlines fail to adopt protective measures. In fact, the decision by all agents to remain unprotected may be a Nash equilibrium, even though from both the vantage points of each individual unit and of society as a whole there are net benefits to everyone from investing in protection. However, in contrast to the prisoner's dilemma problem, there may also be a Nash equilibrium where some or all agents will want to be protected. The challenge is to find ways to convince each of the agents that it is in their best interest to invest in security. Although we initially focus on the airline security case, there are other interesting and topical problems that have similar, though subtly different, analytical structures.

One of these problems concerns the security of a computer network. It is generally the case that once a hacker or virus reaches one computer on a network, the remaining computers can be easily compromised. Because of this possibility of contagion, the incentive that any computer owner has to secure his machine is reduced if he believes that other machines on the network will be insecure. Fire safety in apartment buildings has an analytically similar structure – the risk that an apartment faces depends on the chances of a fire originating in ones own unit as well as the risk of a fire spreading from elsewhere. The incentive an apartment dweller has to take fire precautions, such as installing a sprinkler, depends on her expectations about the policies that will be chosen by other residents in the building. In both of these cases the expectation that others will not adopt protective measures reduces the incentive that a particular agent has to incur these costs. As the number of agents goes to infinity this incentive approaches zero.

To our knowledge this problem of *interdependent security (IDS)* has not been examined in the literature. Orszag and Stiglitz (2002) develop a model for the optimal size of a fire department and point out that homeowners fail to take into account the positive externalities associated with reducing damage to their neighbors by building safer homes. They also note that an increase in government investment in security will tend to reduce individual investment. What they do not show is that the economic incentives for investing in preventive measures decrease as the number of unprotected homes increases, which implies that there is an optimal scale of neighborhood development. There is thus a need for either public sector intervention or coordinating mechanisms to induce this activity and reduce the need for larger fire departments.

One question that the present paper addresses is how to induce tipping mechanisms as characterized by Schelling (1968). In other words, how can one ensure that enough agents will invest in security so that all the others will follow suit? At some level this aspect of the problem is similar to the phenomena that arise with network externalities, where a community will standardize on one of several competing products after enough members have adopted a particular product. (Arthur, 1994; Heal, 1999). In this context the incentive for any agent to invest in security is an increasing function of how many others have already done so. Once a critical mass has invested, then all others will want to do the same.[1]

The next section of the paper develops a model of agents whose security is interdependent by focusing on the airline baggage transfer problem. It illustrates the nature of the externalities that create a disincentive to invest in protection. Section 3 discusses how one can

[1] See Heal (1994, 1999b) for a similar concept of "minimum critical coalition" in the context of interdependency via environmental externalities.

internalize these externalities through different policy tools. Section 4 examines similarities and differences between the airline security case and other IDS problems, namely computer security, fire protection, vaccines, protection against bankruptcy and theft. The concluding section discusses future theoretical and empirical research in this area.

2. A Model of Interdependent Security

Consider a 1 period model where there are *n* risk-neutral agents designated by A_i $i=1...n$. These are the primary actors who have to choose whether or not to invest in security. This choice is taken to be discrete: invest or not invest. In the airline scenario, these are airlines choosing whether or not to invest in a baggage screening system for luggage that is being checked.

Each agent faces the risk of a loss of magnitude *L*. There are two possible ways in which a loss can occur: it can either be initiated on the agent's own property or on the property of another agent. The probability of a loss arising on the agent's own property if it has not invested in security precautions is *p*, so that the expected loss from this event is *pL*. If it has invested in security precautions then this risk is assumed to be zero. The situation is completely symmetric and all agents are identical. For the airline scenario, thorough scanning of baggage that an airline checks on its own will prevent damage from these bags, but there could still be an explosive in a bag transferred from another airline. There is thus an additional risk of loss due to contagion from another agent who has not invested in loss prevention, denoted by *q*.

These probabilities are interpreted as follows. On any given trip there is a probability *p* that an airline without a security system loads a bomb that explodes on one of its own planes.[2] With respect to the chances of contagion, *q* is the likelihood that on any trip a dangerous bag is

[2] All airline trips are assumed to be identical.

loaded onto the plane of one airline and is then transferred to another airline where it explodes.[3] We assume that there is not enough time for an airline to examine the bags from another airline's plane before they are loaded onto its own plane.[4] If there are $n \geq 2$ airlines, the probability per trip that this bag will be transferred from airline i to airline j is $q/(n-1)$. Note that the probability per trip that a bag placed on an airline without a security system will explode in the air is $p+q$.

We assume throughout that the damages that result from multiple security failures are no more severe than those resulting from a single failure. In other words, damages are **not** additive. In the airline baggage scenario, this amounts to an assumption that one act of terrorism is as serious as several. The key issue is whether or not there is a failure, not how many failures there are. Indeed as the probabilities are so low, single occurrences are all that one can reasonably consider. One could think of the definition of a catastrophe as being an event so serious that it is difficult to imagine an alternative event with greater consequences. We focus first on the case of two airlines, each of which is denoted as an agent. This example presents the basic intuitions in a simple framework. We then turn to the multi-agent case.

The 2- Agent Problem

Assume that each agent has perfect information on the risks and costs of protection and has to make a choice between investing in security, **S**, or not to do so, **N.** Think of **S** as investing in baggage screening, and **N** as not doing so. Table 1 shows the payoffs to the agents for the four possible outcomes:

---

[3] The values of p and q are assumed to be exogenous. In the case of terrorism these probabilities may change as a function of the type of security measures undertaken by the airlines. The case of endogenous probabilities is treated in Heal and Kunreuther (2002a).

[4] This is the current practice for all airlines except El Al who does screen bags transferred from other airlines. In fact, the destruction of flight Pan Am 103 in December 1988 over Lockerbie was due to a bomb checked in Malta and then transferred to Pan Am 103 in London via Frankfurt. (http://www.cbc.ca/news/indepth/lockerbie/investigation.html). The transferred piece of luggage was not inspected at either Frankfurt Airport or at Heathrow Airport in London.

Table 1: Expected Outcomes Associated with Investing and Not Investing in Security

		Agent 2 (A_2)	
		S	***N***
Agent 1 (A_1)	***S***	*Y-c, Y-c*	*Y-c-qL, Y -pL*
	N	*Y-pL, Y-c-qL*	*Y –[pL + (1-p) qL], Y –[pL+ (1-p)qL]*

Here Y is the income of each agent before any expenditure on security or any losses from the risks faced. The cost per trip of investing in security is c. The rationale for these payoffs is straightforward. If both invest in security, then each incurs a cost of c and faces no losses so that their net incomes are $Y-c$. If A_1 invests and A_2 does not (top right entry) then A_1 incurs a cost of c and also runs the risk of a loss emanating from A_2. The probability of A_2 contaminating A_1 is q, so that A_1's expected loss from a bomb originating elsewhere is qL. This cost represents the negative externality imposed by A_2 on A_1. A_2 incurs no baggage security costs and faces no risk of contagion from A_1, but it does face the risk of losses originating at home, pL. The lower left payoffs are just the mirror image of these. If neither agent invests in security, then both have an expected payoff of $Y- pL - (1-p) qL$.

Now that the outcomes have been specified, one can ask the natural question: under what conditions will the agents invest in security? It is clear from Table 1 that for investment in security to be a dominant strategy, we need

$$Y-c>Y-pL \quad \text{and } Y-c-qL>Y-pL-(1-p)\, qL$$

The first inequality just says that $c<pL$: the cost of investing in security must be less than the expected loss, a natural condition for an isolated agent. The second inequality is more interesting: it reduces to $c<pL- pqL = pL(1-q)$. This is clearly a tighter inequality reflecting the possibility of contagion from the second agent. This possibility reduces the incentive to invest in security. Why? Because in isolation, investment in security buys the agent complete freedom

from risk; with the possibility of contagion it does not. Even after investment there remains a risk of loss emanating from the other agent. Investing in security buys you less when there is the possibility of contagion from others.

In the 2-agent problem with identical costs, one can determine the optimal behavior of each agent if they both make decisions simultaneously without any communication. In this non-cooperative environment if $c < pL(1-q)$, then both agents will want to invest in protective measures **(S,S)**; if $c > pL$ then neither agent will want to invest in protection **(N,N)**. If $pL<c<pL(1-q)$ then there are two Nash equilibria **(S,S)** and **(N,N)** and the solution to this game is indeterminate.[5]

If the agents have different costs of investing in security measures, then there may be a Nash equilibrium where one agent invests in security and the other does not. Specifically, let c_1 and c_2 be the costs of the two agents: then (**N,S**) will be a Nash equilibrium if $c_1>pL$ and $c_2<pL(1-q)$. This mixed equilibrium requires that the two costs differ by at least pqL.[6]

The solution concept for two agents with identical costs and risks is illustrated below with a numerical example. Suppose that $p= .2$, $q=.1$, $L=1000$ and c= 185. The matrix in Table 1 is now represented as Table 2.

Table 2: Expected Costs Associated with Investing and Not Investing in Security for Illustrative Example

		Agent 2 (A_2) ***S***	***N***
Agent 1 (A_1)	***S***	*Y-185, Y-185*	*Y-285, Y-200*
	N	*Y-200, Y-285*	*Y-280, Y-280*

[5] We have excluded the cases where the agent is indifferent between either investing or not investing in security (i.e.,where $c=pL$ or $c=p(1-q)L$.

[6] See Heal and Kunreuther (2002a) for a treatment of the case where agents have heterogeneous risks and/or costs of protection.

If A_2 invests in security (**S**), then it is worthwhile for A_1 to also invest in security, since without protection its expected losses will be *pL= 200* and it will only have to spend 185 to eliminate this risk. If A_2 does not invest in security (**N**), then there is still a chance that A_1 will experience a loss even if it protects itself. The expected benefits to A_1 of investing in security will now only be $pL(1-q) = 180$, which is less than the cost of the security measure. Hence A_1 will **not** want to invest in protection. In other words, either both agents invest in security or neither of them do so. These are the two Nash equilibria.

Multiple agents

Let us now turn to the more general case of *n* identical agents all symmetrically placed. If all but one of the agents have invested in security, then the risk facing the remaining one is identical to what it would be in isolation: there is no risk of contagion.[7] At the other extreme, suppose none of the other $n-1$ agents have invested in security; then if the remaining agent is protected it still faces risks originating at $n-1$ other locations.

Consider the case of three agents, denoted A_i $i=1,2,3$. Interpret them for concreteness as airlines deciding whether or not to install baggage security systems when bags are only transferred once between airlines. In how many ways can airline 1 (A_1) be victim of a bomb attack if it has instituted a baggage security system but none of the other airlines have done so? Airline 1 can suffer damage from a bag checked onto A_2 and then transferred to A_1. This event occurs with probability $q/2$ since we assume that the bag from airline 2 has an equal chance of being transferred to either A_1 or A_3. A bomb-laden bag from A_3 can also damage A_1. This can occur when A_2 does **not** transfer a dangerous bag to A_1 but A_3 does. This event occurs with

[7] When all other agents have invested in protection the remaining agent has a type of herd immunity, a term used with respect to the spreading of diseases. With respect to contagion from diseases there is no incentive for an individual to be vaccinated if everyone has been, and the disease can only be contracted from another individual who has it. In the airline security case there is a greater incentive to invest in protection if there is no chance of contagion from others. The similarities and differences between airline security and vaccinations are discussed in Section 4.

probability $(1-q/2)q/2$. Since all agents are assumed to be identical, the negative externalities are the same for every airline.

Define $X(3,0)$ as the expected negative externality to any airline i that has installed a baggage checking system if there are 3 airlines and none of the others have instituted this security measure. $X(3,0)$ is given by $(q/2)\,[1+(1-q/2)]L$. When one other airline has installed a security measure, then the expected negative externality $X(3,1)$ is given by $(q/2)L$, since there is only one airline without a security system and it transfers a contaminated bag to airline i with probability $q/2$. .

If there are four airlines then the expected negative externalities become:

$X(4,2) = (q/3)\,L$
$X(4,1) = (q/3)\,[1+(1-q/3)]\,L$
$X(4,0) = (q/3)\,[1+(1-q/3)+(1-q/3)^2]L$

For $n>1$ agents this generalizes to

$$X(n,0) = [q/(n-1)] \sum_{t=0}^{n-2} [\,[1-q/(n-1)]^t]L = \{1-[1-q/(n-1)]^{n-1}\}\,L \quad (1)$$

The limit of this expression as n tends to infinity is

$$\lim_{n\to\infty} X(n,0) = (1-e^{-q})\,L$$

We can summarize this in the following proposition:

Proposition 1: If there are n identical agents, none of whom has invested in security, then the expected loss inflicted on any agent by all others is $X(n,0)= \{1-[1-q/(n-1)]^{n-1}\}\,L$. As $n\to\infty$, this converges to $(1-e^{-q})\,L$.

If $q=0$ then $X(n,0)$ is zero for any n; When $n=\infty$, $X(\infty,0)$ increases monotonically in q reaching its largest value of $0.63L$ when $q=1$. Intuitively this makes sense. With a zero chance of baggage transfer there is no negative externality. If bags with explosives are transferred to other airlines with certainty, then in the limit the expected negative externality to any airline is *63%* of the possible loss. For a given value of q, the term $X(n,0)$ decreases monotonically as n increases, taking on the value of qL for $n=2$ and falling to $(1-e^{-q})L$ as n approaches infinity. So the externality is largest when there is only one other airline and decreases as the number of airlines rises. Again there is an obvious intuition here: as the number of airlines increases, the chance of a transferred bag reaching any particular airline on any given trip falls.

When there are n firms, the payoff to A_1 from not investing in security when the other $n-1$ are also not investing is

$$Y - pL - (1-p)X(n,0) \qquad (2)$$

The payoff to A_1 from investing is

$$Y-c - X(n,0) \qquad (3)$$

Comparing *(2)* and *(3),* investing is the better strategy if and only if

$$c < p[L - X(n,0)] \qquad (4)$$

Equation *(4)* implies that there is less incentive to invest in protection with higher negative externalities associated with contagion.

What is the structure of the set of possible Nash equilibria? For the two-agent case, **(S,S)** is a dominant strategy equilibrium if $c<pL(1-q)$ and a Nash equilibrium if $c<pL$. The strategy **(N,N)** is a dominant strategy equilibrium if $c>pL$ and a Nash equilibrium if $c>pL(1-q)$. There is an interval $pL(1-q)<c<pL$ in which both **(S,S)** and **(N,N)** are Nash equilibria. For the n-agent case **(S,S…S)** is a dominant strategy equilibrium if $c < p[L-X(n,0)]$ and **(N,N….N)** is a dominant strategy equilibrium if $c > pL$. When c is between these two values there are two stable Nash equilibria---(**N,N…N)** and **(S,S….S).**

Could there be other Nash equilibria? The answer is **no** if all agents are identical. In the two agent case for **(N,S)** to be an equilibrium it is necessary that $Y-pL>Y-c$ or $c>pL$ and also that

$Y\text{-}c\text{-}qL > Y\text{-}pL\text{-}(1\text{-}p)qL$ or $c < pL(1\text{-}q)$ which is obviously impossible. So the only Nash equilibria are where both agents invest or both do not invest. This does not change as the number of agents increases. Even with many identical agents, they all will choose the same strategy.

Proposition 2: Consider a Nash equilibrium in the *n*-agent problem (n>2) defined above. Each agent has two strategies **N** and **S** with payoffs described by equation *(2)* if the agent does not invest in security and equation *(3)* if it does invest in security. The only Nash equilibria are ones where all agents choose the same strategy.

***Proof*.** Assume to the contrary that (**S,S,…..S,N**) is a Nash equilibrium. Without loss of generality we assume the last agent chooses **N** and all others choose **S**. Then for all agents from *1* to *n-1* this implies that the strategy **S** is a best response to *n-2* agents choosing **S** and one choosing **N**. But for agent *n*, **N** is the best response if there are *n-1* agents choosing **S.** However if **S** is a best response to *n-2* agents choosing **S** and *1* agent choosing **N**, then it is also a best response to *n-1* agents choosing **S**; changing one choice from **N** to **S** will reduce the externality on all other agents and increase the incentive to choose **S**. This contradicts the assumption that agents choose differently at a Nash equilibrium. The same argument can be applied to cases in which more than one agent chooses a divergent strategy.

There are three critical values of *c* that need to be considered in determining the nature of the equilibria when there are *n* identical agents in the system. Let c^{**} represent the value of *c* above which an agent will **not** invest in protection in isolation. Clearly $c^{**}=pL$. Let $c^{*}(n, 0)$ represent the value of *c* below which an agent will still want to invest in security even if none of the other *n-1* agents are protected. Then $c^{*}(n, 0) = p[L - X(n,0)]$. For the above example where $L=1000$, $p=.2$ and $q=.1$ and $n=10$ airlines, then $X(10,0)=19.1$, $c^{**}=200$ and $c^{*}(10,0) = 180.9$. When $c>200$ then the only Nash equilibrium is (**N,N,...N**) and none of the agents will want to invest in protection. If $c < 180.9$ then the only Nash equilibrium is (**S,S,....S**). For $180.9 < c < 200$ there are two Nash equilibria (**N,N,....N**) and (**S,S,.....S**).

The cost of investing in protection plays a key role in determining the nature of the Nash equilibrium. For sufficiently low values of c ($c < c^*(10,0) = 180.9$ in this example), agents will want to invest in security even if they can be harmed by others because they are able to reduce the expected losses from a loss originating on their own property sufficiently to justify protecting themselves. On the other hand, if c is sufficiently high ($c > c^{**} = 200$ in this example) then it is not worthwhile for any agent to protect itself against a loss originating on its own property even if it knows it will not suffer contagion from others. When the cost of protection is between these two values then the decision to invest in protection is influenced by what the other agent does.

Relationship Between IDS and Prisoner's Dilemma Problems

The problem of encouraging individuals to adopt protective measures resembles the prisoner's dilemma problem in the sense that it is often advantageous for all agents to adopt protection for both themselves and society, but none of them have an economic incentive to do so on their own**.** A classic prisoner's dilemma is where each firm has a cost incentive to undertake some activity (e.g. polluting the environment). It knows that if there were a coordinating mechanism so that none of them engaged in this activity, they would each be better off and social welfare would also be improved. (e.g., each firm's profits would be higher and the air and/or water would be cleaner.)

For certain cost structures the IDS problem has the same characteristics as a prisoner's dilemma. More specifically in the 2-agent case if $pL < c$ then each agent prefers not to invest in security [i.e. (**N,S**) > (**S,S**) for agent 1 and (**S,N**) > (**S,S**) for agent 2] leading to a single Nash equilibrium at (**N,N**) . However, if $pL + (1-p)qL > c$ then both agents would be better off at (**S,S**) than at (**N,N**).

For other situations where $pL > c > pL(1\text{-}q)$, the IDS problem differs from the prisoner's dilemma, since there are two Nash equilibria (**S,S**) and (**N,N**). Now, for example, if agent 1 can convince agent 2 to invest in security, there will be an economic incentive for agent 1 to voluntarily follow suit; otherwise both agents will be content to not invest in protection knowing that they would both be better off if they each undertook protective measures. For these parameter values we have a coordination problem [Heller (1986), Crawford and Haller (1990) and Van Huyck et al. (1990)] Coordination problems arise in systems that may have multiple Pareto ranked equilibria, where all agents have an interest in moving to the highest-ranked equilibrium but there is no institutional mechanism to attain this outcome.[8] In the context of the airline example, an airline is more likely to invest in a baggage security system if it knows that the other airline has taken this step.

3. **Internalizing Externalities**

One way to encourage agents to invest in security when they face the possibility of contagion from others is to internalize the externalities. This section examines the roles that different policy tools ranging from private market mechanisms to government regulations to collective choice can play in encouraging agents to adopt protective measures for IDS problems.

Insurance

Insurance discourages investment in security if insurers face moral hazard problems due to their inability to detect careless on the part of the insured agents who know that they will receive compensation should they suffer a loss. In this case one may actually lose an **(S,S,….,S)** equilibrium if the parties are allowed to insurance themselves against losses.[9] If moral hazard problems can be eliminated through the terms of the contract (e.g. deductibles, coinsurance)

[8] See Chapter 7 in Camerer (2003) for a comprehensive summary and analysis of how players make choices in coordination games in controlled laboratory experiments.

[9]We thank Richard Zeckhauser for pointing this out to us.

and/or through monitoring and inspection, then insurance with actuarially fair premiums encourages a risk averse individual or firm operating in isolation to adopt protection whenever the cost of the measure is less than the reduction in the expected losses.

To deal with the externalities created by others who do not invest in security, the unit causing the damage must be forced to pay for the losses. This means that if a bag transferred from Airline 1 to Airline 2 were to explode, then Airline 1's insurer would be required to pay for the damage to 2. This is not how current insurance practice operates. An insurer who provides protection to A_i is responsible for losses incurred by agent i no matter who caused the damage.[10] One reason for this contractual arrangement between insurer and insured is the difficulty in assigning causality for a particular event.[11] A single insurance program that provided coverage to all agents would, however, want to internalize the externality.

To illustrate this point consider the IDS case with two identical agents (A_1 and A_2). Suppose each agent had its own insurer who charged a premium based on expected losses. A_1 contacts its insurer inquiring about a premium reduction for undertaking a protective measure, knowing that $c < pL$. If the insurer knows or suspects that A_2 has not invested in protection, it will only be willing to reduce the premium by $p(1-q)c$ because of the contagion effects from A_2 to A_1. On the other hand, a single insurer covering both agents, that is a monopolist or a social insurance program, can require both A_1 and A_2 to invest in the protective measure and in return give each agent a premium reduction of pL.

Liability

[10] If the damage from an insured risk is due to negligence or intentional behavior, then there are normally clauses in the insurance policy that indicate that losses are not covered (e.g. a fire caused by arson).
[11] With respect to fire damage, a classic case is H.R. Moch Co., Inc. v Rensselaer Water Co. 247N.Y.160, 159 N.E. 896 which ruled that "A wrongdoer who by negligence sets fire to a building is liable in damages to the owner where the fire has its origin, but not to other owners who are injured when it spreads". We are indebted to Victor Goldberg who provided us with this case.

If an agent who caused damage to other agents by not adopting a protective measure were held liable for these losses, then the legal system would internalize the externalities due to interdependent security. For the two-agent example, suppose that A_1 knew that by not investing in security it would be liable for damage that it caused to A_2. It would then invest in security whenever $c < (p+q)L$.

Although the liability system has attractive theoretical properties, it faces practical problems due to high transaction costs. Determining the cause of the loss can be very costly and extremely time consuming. In the case of the airline example, it would be difficult to know whether an unchecked bag from another airline caused damage to the plane or whether it was due to one of the airline's own bags. For example, in the PanAm 103 case it took many months of expert forensic work to determine what bag caused the crash and where it came from (Lockerbie Verdict 2001). The costs of settling these disputes appear to favor a liability system where each agent is responsible for its own losses unless there is a clear case of negligence.

Fines and Subsidies

The public sector could intervene directly in IDS problems by levying a fine of *F* on any entity that does not invest in security, or alternatively providing an entity with a subsidy of *G* to encourage protection. Consider the case of fines. With identical agents one would want the fine to be high enough so that the only Nash equilibrium would be (**S,S....S**). The magnitude of *F* depends on the number of agents and the cost of protection, *c*.

Suppose that there are *n* agents in the pool and none of them have invested in security. The government wants to determine the minimal fine *F** to induce each agent to protect itself. As shown in section 2 the costs to an agent who invests in protection will be

$$c + X(n,0).$$

If an agent does not invest in protection and is fined *F* dollars, its cost will be

$$pL+(1-p)X(n,0)+F$$

Hence for any agent to want to invest in protection when no one else does, the fine must be high enough so that

$$F > c- p[L- X(n,0)]$$

If $c < p(L- X(n,0))$ then there is no need to impose any fine on an agent for it to want to invest in protection. Hence

$$F^* > \max \{0, c-p(L- X(n,0)] \}$$

Consider the airline example where $n=10$ and $X(10,0)= 19.1$ with $F=0$. An agent will only invest in security if noone else does if $c< p[L-X(10,0)]=80.9$. If $c> 80.9$, then $F^*= c – 80.9$.[12]

A subsidy G for adopting protective measures plays the identical role in inducing agents to invest in security as a fine with one major difference: the subsidy has to be paid to induce the agent to invest in security while a fine will not be incurred by the agent if it adopts the appropriate protective measure.[13] G reduces the cost c to the agent, thus making the protective measure more attractive. In the above example, if $c < 80.9$ no subsidy will be necessary to induce an agent to invest in security. Otherwise, the minimal subsidy $G^* = c - 80.9$.

Regulations and Third Party Inspections

The possibility of contagion from other units provides a rationale for well-enforced regulations that require individuals and firms to adopt cost-effective protective mechanisms when they would not do so voluntarily. In the identical n-agent example, a regulation would be viewed as desirable from both private and social welfare perspectives under the following conditions:

---

[12] Suppose that $c>c^{**}$ so that there is no incentive for any agent to invest in protection even if all other $n-1$ agents have protected themselves. If there are additional indirect benefits from protection besides a reduction in the expected loss (pL), then the government may want to impose a fine on unprotected agents that is high enough to induce each of them to protect itself.

[13] We thank Stan Baiman for pointing this out to us.

- there are two stable Nash equilibria **(S,S,….S)** and **(N,N….N)**
- the equilibrium **(S,S…S)** yields higher profits for all agents than **(N, N…N)**
- none of the agents voluntarily adopted protective measures because they believed others would not do so.

One would thus want to consider a regulation whenever $p[L-X(n,0)]<c< pL$. Each agent would be better off if it was forced to invest in security, knowing that all the other agents were required to do the same. In this case regulation solves the coordination problem. There may also be a need for well-enforced regulations if there were externalities to other parties in addition to the contagion effects between the agents. For example, when a building collapses it may create externalities in the form of economic dislocations and other social costs that are beyond the economic losses suffered by the owners. These may not be taken into account when the owners or developers evaluate the importance of adopting a specific mitigation measure and hence may justify the need for building codes. [Cohen and Noll 1981; Kleindorfer and Kunreuther (1999)].

One way for the government to enforce its regulations is to turn to the private sector for assistance. More specifically, third party inspections coupled with insurance protection can encourage divisions in firms to reduce their risks from accidents and disasters. Such a management-based regulatory strategy shifts the locus of decision-making from the regulator to firms who are now required to do their own planning as to how they will meet a set of standards or regulations. [Coglianese and Lazer (2001)].[14]

Coordinating Mechanisms

Rather than relying on government regulations, one could turn to the private sector to coordinate decisions through industry associations. In the context of the illustrative example of

[14] Kunreuther, McNulty and Kang (2002) show more formally how such a program could be implemented in practice.

airline security in Section 2, the International Air Transport Association (IATA), the official airline association, could have made the case to all the airlines that they would be better off if each one of them utilized internal baggage checking so that the government would not have had to require them to so.[15]

An association can play a coordinating role by stipulating that any member has to follow certain rules and regulations, including the adoption of security measures, and has the right of refusal should they be asked to do business with an agent that is not a member of the association and/or has not subscribed to the ruling. IATA could require all bags to be reviewed carefully and each airline could indicate that it would not accept in-transit bags from airlines that did not adhere to this regulation. By receiving a seal of approval from IATA, the airline would also increase its business since passengers would shun airlines that were not part of the agreement.[16]

Another solution would be for airlines that had invested in security to announce publicly that they will not accept passengers and hence baggage from any airline that doesn't have security. They would then public announce to all prospective passengers which airlines fell in this category.[17] This tactic may encourage the unprotected airlines to invest in security because of their fear of losing customers in the future.[18]

4. Protective Measures for Other IDS Problems

Each IDS problem has its own structure that calls for certain levels of protection. This section illustrates similarities and differences between the airline security and other IDS

---

[15] If all the airlines felt that they could not afford to do this even if everyone adopted these measures, then they would resist any attempt by IATA to require them to take this step and request the government to pay for these security measures, as they have done.

[16] IATA follows this type of policy in agreements regarding to transfer of tickets. An IATA-affiliated airline will not honor a non-IATA airline ticket unless it conforms to the IATA tariff conference. See the IATA web site at http://www.iata.org/membership/steps.asp#10.

[17] We thank Jack Hershey for suggesting this option to us.

[18] On a more informal level it might be possible to establish social norms that generate pressure to invest in protection. See Sunstein (1996) for a more detailed discussion of social norms. Ostrom (1990 Chapter 6) deals with the conditions under which norms evolve governing the use of common property resources.

problems. For each context, the definitions of p and q depend on the way losses directly impact an agent and the nature of the process of contagion. These characteristics affect the nature of the Nash equilibria, and hence the optimal strategy for improving private and social welfare.

Computer Security

Protecting computer networks from viruses and from hackers reduces the chances that a loss will occur to the agent who takes protection while at the same time reducing negative externalities. Each agent on the network can make its own investment in protection against external attacks where it would be the target, but not against a virus that would come from the internal network, which is supposed to provide a "friendly" source of information. In other words, the effectiveness of this investment depends on those made by others. If one computer is unprotected then malicious external agents could attack the entire system via this one computer. A hacker who gains access to a network via one weak link can in many cases compromise all computers on the network through internal access. One unprotected node can endanger all the other nodes in an interconnected network even if they have invested in protection against direct external attacks.[19] In the airline security problem the configuration is different as only one plane can be affected by a contaminated piece of luggage---a bomb can only explode once.

The definitions of p and q also differ in the computer network example from the airline security case. Now p is the probability that a computer is infected with a virus and q is the probability that a computer with a virus contaminates other computers in the system. In other words, $q \leq p$. As one unprotected computer can impact internally all $n-1$ other computers whether or not they are protected, the expected negative externalities associated with the

[19] We are indebted to Yechiam Yemini for this information. See also Anderson (2001) for a discussion of the types of incentives that can be utilized to deal with information security and the need for engineers, economists, lawyers and policymakers to join forces in dealing with this problem.

computer security problem are much greater than for airline security. More specifically if computer i is protected against viruses and all the other n-1 computers are not, then

$$X(n,0) = qL \sum_{t=0}^{n-2} [(1-q)^t] = [1-(1-q)^{n-1}]L$$

This implies that as n increases, $X(n,0)$ also increases so that there is less incentive for any agent to invest in protecting its computer system due to the increased chance of contagion from others. As the number of agents increases without limit then

$$\lim_{n\to\infty} X(n,0) = qL \sum_{t=0}^{\infty} [(1-q)^t] = qL/[1-(1-q)] = L \qquad (5)$$

This implies that in the limit $c < p[L-X(n,0)] = 0$ so that investing in computer security can never be a dominant strategy as long as the cost of protection is positive.

A comparison of this result with *Proposition 1* for the airline case is instructive. When $n=2$ the two cases are by definition the same. For $n = \infty$ and $q=1$ the airline negative externality is $0.63L$, whereas in the computer network case the number is just L. A computer virus is a public bad – it's capacity to damage is non-rival – whereas a bomb on an airline is a private bad.

Fire Protection

Investing in sprinkler systems in an apartment in a multi-unit building to reduce the potential losses from fire has a problem structure similar to the computer security problem. A fire that starts in an unprotected apartment can spread to other units and damage them whether or not they have sprinkler systems installed. If a fire in any unit could spread to all the other units simultaneously, then this problem would be identical in structure to a computer virus. In reality a fire normally destroys units only on the same or adjacent floors of buildings, in which case any apartment unit would only be subject to damage from at most *m* of the *n* units in the building. [20] Suppose apartment 1 was protected with a sprinkler system and all the other *m* units surrounding it were unprotected. Then the negative externalities it would incur would have an expected value given by

$$X(n,0) = qL \sum_{t=0}^{m-1} [(1-q)^t]$$

As the number of units *m* that can impact on a given apartment decreases, then *X(n,0)* also decreases and the apartment unit will be more likely to invest in protection.

Vaccinations

The decision on whether to get vaccinated has the following feature that makes it an IDS problem: if I am vaccinated against a contagious disease, you will not catch it from me. So one person investing in protection conveys positive externalities on others, as in the airline security, computer security and fire protection problems.[21]. Consider the Nash equilibrium that may arise when people decide whether or not to be vaccinated. Suppose that tomorrow an effective vaccine against influenza is approved for general use. When choosing whether to be vaccinated or not, each person has to anticipate the choices of others. If everyone else were to be vaccinated, then there would be no point in my being vaccinated, as I would be in no danger of catching the flu, unless I could get it from an external source. At the other extreme, if I believed that most people would not be vaccinated, this would increase my incentive to be vaccinated.

[20] One of us, Heal, lives in an apartment building and was recently told by the building's insurance agent that a serious fire usually destroys the floor on which it starts, and normally damages two floors above via smoke and flames and two below through water damage from extinguishing the fire.
[21] Philipson (2001) has a nice summary of recent research on economic epidemiology and the role that vaccines play in reducing the spread of diseases.

From this we can see that if the vaccination cost is sufficiently low and the risk is sufficiently high then the situation where no one is vaccinated cannot be a Nash equilibrium.[22] On the other hand, everyone being vaccinated is also not a Nash equilibrium if one can only get an illness or disease from someone who already has contracted it and a vaccination provides complete protection. The Nash equilibrium will now be a mixture of **N**s and **S**s even when all individuals are identical. Some individuals will decide to get vaccinated while others prefer to be unprotected. Those who decide to get vaccinated will have no incentive to change their minds because there will be enough people who are unprotected, so that the chances of contracting the disease will be greater than the expected cost and potential side-effects of the vaccine. Similarly those who have not protected themselves will find that the expected costs and side effects of the vaccine will exceed the expected benefits from being protected.[23]

Protection Against Bankruptcy

Another example of an IDS problem is the decision by a unit in a multi-divisional firm as to whether or not to invest in protective measures that reduce the chance of the firm going bankrupt. The economic incentive for any division in a firm to invest in risk-reduction measures depends on how it expects the other divisions to behave in this respect. Consider Division 1. If it thinks that the other divisions will not invest in protection, then this reduces Division 1's incentive to do so. On the other hand should Division 1 believe that the others are taking appropriate steps to mitigate their risks, then it may be best for Division 1 to follow suit.

Recently Arthur Anderson was sent into bankruptcy in large part because of the actions of its Houston branch. Several years ago Barings was destroyed by the actions of a single trader

[22] See Hershey et al (1994) for a more detailed discussion of the role that free riding, altruism, and bandwagoning play in vaccination decisions.

[23] We thank Richard Zeckhauser for pointing out this feature of the Nash equilibrium to us. Heal and Kunreuther (2002b) derive Nash equilibria for the vaccination problem when there is the possibility of contracting the disease from an outside source and there are indirect externalities associated with transferring the disease from others.

in its Singapore branch. In each of these cases the risk of bankruptcy faced by any unit was affected by its own choices and by those made by others in the firm. A culture of risk-taking can spread through the organization because the knowledge that a few groups are taking risks reduces the incentives that others have to manage them carefully.

One of the major purposes of firms is to internalize these types of externalities. However, it may be difficult for them to deal with this problem if they have a decentralized organizational structure. Can they encourage their divisions to invest in risk-reducing measures through internal rules? Can they establish a culture of safety so that each division feels obligated to invest in protective measures because it reduces the potential for catastrophic losses? The answers to these questions are not obvious when there is the possibility of any division bringing down the entire firm and the costs of taking preventive action can be costly.[24]

Theft Protection

Consider the case where a burglar is considering which one of a set of identical houses in a neighborhood to rob. One of his principal concerns is the likelihood of being caught when attempting to break into the house. By installing a burglar alarm you increase the chances that the intruder will be detected. If you announce publicly with a sign that your house has been protected, then the burglar will often look for greener pastures to invade. In other words, installing a burglar alarm in your house, and announcing it, decreases the chances that your house will be robbed and increases the likelihood that other unprotected homes will be targets for the burglar.[25]

Let p be the probability of a loss (L) to any house when none of the homes in the area have invested in protection. For example, if a thief randomly chooses one of the n houses in the

[24] The challenges associated with protecting a multi-divisional firm from bankruptcy are discussed in Kunreuther and Heal (2002).
[25] We appreciate a helpful discussion with Daniel Kahneman on this point

area as a target, then $p = 1/n$. Now suppose that you purchase a burglar alarm that can always detect a thief should he attempt to break into your house and you publicize that your house is protected in this way.[26] The risk of a loss to your house is now zero, independent of what other houses have done. If you protect yourself against theft, there is now an increase in the probability that one of the other nearby houses will be robbed. Let p' represent this revised probability of a theft with $p' > p$. In the case of random theft, your house is off-limits and the other $n-1$ houses have a $p' = 1/n-1$ chance of being burglarized.

If the cost of the burglar alarm is c, all houses are identical and individuals are risk neutral, then no one will invest in a burglar alarm if $c > pL$. If $c < pL$ then everyone will want to protect themselves. Note that the Nash equilibrium is a static concept. If for whatever reason some individuals invest in a burglar alarm (e.g. they are required to do so by their insurance company), then others will have an increased incentive to also invest since their chances of being burglarized has increased.

Suppose that instead of publicly revealing that one has a burglar alarm, you and others in the neighborhood connect the alarm to the local police station, so that criminals cannot determine who is protected. In this case installing an alarm system does not reduce the probability that an individual house will be broken into. It may, however, provide a positive externality by reducing the chances of a crime occurring in your community if the burglars know that a certain percentage of homes are protected. This type of unobservable precaution is similar to the Lojack car retrieval system. Ayres and Levitt (1998) show that the marginal social benefit of an additional unit of Lojack is as much as 15 times greater than the marginal social cost in high

[26] Richard Zeckhauser tells the story of installing a phony burglar alarm many years ago which consisted of a panel with little red light on the outside of his house. He now has a real burglar alarm system but believes the panel is still effective since he has not any attempted thefts since the real alarm was installed.

crime areas. However, those who install Lojack in their cars obtain less then 10 percent of the total social benefits associated with this protective measure.

5. Future Research

Deciding whether to invest in risk reduction where there is some interdependency between your actions and those of others raises a number of interesting theoretical and empirical questions. This paper has examined the case in which all agents are identical. Heal and Kunreuther (2002a) consider situations where the agents have differential protection costs and risks, and where the actions of those creating potential losses are impacted by agents' protective decisions, as in the case of terrorism. Open questions are how agents behave in multi-period models and what are the appropriate behavioral models of choice for characterizing individuals who make imperfectly rational decisions.

The issues discussed above also suggest a number of empirical studies on interdependent security. Given the concern with terrorism both in the United States and the rest of the world it would be interesting to learn more about what factors lead some organizations to invest in security and why others are deterred from doing so. What institutional mechanisms would aid the decision process of agents regarding protective measures when others will be affected? Can industry associations, like IATA for the airlines, play an important role in facilitating actions by individual companies? What are the appropriate roles of the public and private sectors in developing strategies that include economic incentives (fines or subsidies), third party inspections, insurance coupled with well-enforced regulations and standards? These are natural extensions of this problems and topics that deserve future research.

REFERENCES

Allen, Franklin and Gale, Douglas (2000) “Financial Contagion” *The Journal of Political Economy*, 108:1-33.

Anderson, Ross (2001) “Why Information Security is Hard - an Economic Perspective” Working Paper, Computer Laboratory, Cambridge. http://www.ftp.cl.cam.ac.uk/ftp/users/rja14/econ.pdf.

Arthur, Brian (1994) *Increasing Returns and Path Dependence in the Economy.* Ann Arbor: University of Michigan Press.

Ayres, Ian and Levitt, Steven (1998) “Measuring the Positive Externalities from Unobservable Victim Precaution: An Empirical Analysis of Lojack” *Quarterly Journal of Economics* 113:43-77.

Bolton, Patrick and Farrell, Joseph (1990) “Decentralization, Duplication, and Delay” *The Journal of Political Economy*, 98: 803-826.

Camerer, Colin (2003) *Behavioral Game Theory : Experiments on Strategic Interaction* Princeton, NJ: Princeton University Press

Coglianese, Cary and Lazer, David. (2001) “Management-Based Regulation: Using Private Sector Management to Achieve Public Goals” Regulatory Policy Program Working Paper RWP01-047 Cambridge, MA: John F. Kennedy School of Government, Harvard University

Cohen, Linda and Noll, Roger (1981) "The Economics of Building Codes to Resist Seismic Shocks" *Public Policy* Winter 1-29.

Crawford, Vincent and Haller, Hans (1990) “Learning How to Cooperate: Optimal Play in Repeated Coordination Games” *Econometrica* 58: 571-595.

Heal, Geoffrey (1994). “Formation of International Environmental Agreements,” in C. Carraro (ed.) *The International Dimension of Environmental Policy*, Boston: Kluwer.

Heal, Geoffery (1999a) “Price and Market Share Dynamics in Network Industries”. Chapter 10 of G. Chichilnisky (ed) *Markets, Information and Uncertainty: Essays in Honor of Kenneth J. Arrow.* New York: Cambridge University Press.

Heal, Geoffrey (1999b) “New Strategies for the Provision of Global Public Goods: Learning from International Environmental Challenges,” in I. Kaul and M. Stern (ed) *Global Public Goods* New York: Cambridge University Press.

Heal, Geoffrey and Kunreuther, Howard (2002a) “You Can Only Die Once” (in preparation)

Heal, Geoffrey and Kunreuther, Howard (2002b) “The Vaccination Game” (in preparation)

Heller, Walter (1986) “Coordination Failure in Complete Markets with Applications to Effective Demand” in *Equilibirium Analysis: Essays in Honor of Kenneth J. Arrow Vol II,* ed. W.P.Heller, R.M. Starr and D.A. Starrett, Cambridge University Press.

Hershey, John, Asch David, Thumasathit T, Meszaros Jacqueline, and Waters V, (1994) "The Roles of Altruism, Free Riding, and Bandwagoning in Vaccination Decisions," *Organizational Behavior and Human Decision Processes* 59:177-187.

Kleindorfer, Paul and Kunreuther, Howard. (1999). “The Complementary Roles Of Mitigation And Insurance In Managing Catastrophic Risks." *Risk Analysis, 19:727-38.*

Kunreuther, Howard, McNulty, Patrick and Kang, Yong (2002) “Improving Environmental Safety Through Third Party Inspection” *Risk Analysis.* 22: 309-18.

Kunreuther, Howard and Heal, Geoffrey (2002) “A Firm Can Only Go Bankrupt Once” Wharton Risk Management and Decision Processes Center Working Paper Philadelphia: University of Pennsylvania

Lockerbie Verdict (2001) *Her Majesty’s Advocate, v. Abdelbaset Ali Mohmed Al Megrahi and Al Amin Khalifa Fhimah ,* Case No: 1475/99 (High Court of Justiciary at Camp Zeist— January 31)

Musumeci, James J. and Sinkey, Joseph F. Jr. (1990) “The International Debt Crisis, Investor Contagion, and Bank Security Returns in 1987: The Brazilian Experience” *Journal of Money, Credit and Banking*, 22: 209-220.

Orszag, Peter and Stiglitz, Joseph (2002) “Optimal Fire Departments: Evaluating Public Policy in the Face of Externalities” Brookings Working Paper January.

Ostrom, Elinor. (1990) *Governing the Commons: The Evolution of Institutions for Collective Action.* Cambridge, UK: Cambridge University Press.

Philipson, Tomas (2001) “Economic Epidemiology and Infectious Diseases” in Joseph Newhouse and Anthony Culyer eds. *The Handbook of Health Economics* North Holland Press.

Polonchek, John and Miller, Ronald K. (1999) “Contagion Effects in the Insurance Industry” *The Journal of Risk and Insurance*, 66: 459-475.

Schelling, Thomas (1978) *Micromotives and Macrobehavior* New York: Norton.

Sunstein, Cass (1996) “Social Norms and Social Roles,” *Columbia Law Review* 96:903-68.

Van Huyck John B., Battalio, Raymond C. and Beil, Richard O. (1990) “Tacit Coordination Games, Strategic Uncertainty, and Coordination Failure” *The American Economic Review*, 80: 234-248.