

You Only Die Once: Managing Discrete Interdependent Risks

Geoffrey Heal and Howard Kunreuther^{1 2}

August 2003

¹Heal is at the Graduate School of Business and the School of International and Public Affairs at Columbia University - gmh1@columbia.edu and www.gsb.columbia.edu/faculty/gheal. Kunreuther is the Cecilia Yen Koo Professor at the Wharton School University of Pennsylvania and co-director of the Wharton Risk Management and Decision Processes Center - kunreuther@wharton.upenn.edu. He also is a Visiting Scholar at the Earth Institute at Columbia University.

²We are grateful to Charles Calomiris, Avinash Dixit, Michael Kearns, Nat Keohane, Paul Kleindorfer, Peter Orszag, Todd Sandler, Joe Stiglitz, Thomas Weber and Richard Zeckhauser for valuable comments on various drafts of this paper. We have also benefitted from the comments of participants in NBER Insurance Workshops, the AEA annual meetings in 2003, and seminars at Columbia, Stanford, University of Pennsylvania, Cornell, Ohio State, the London School of Economics and the U.S. General Accounting Office. Responsibility for errors is of course our own. We acknowledge financial support from Radiant Trust and Lockheed Martin.

Abstract

This paper extends our earlier analysis of interdependent security issues to a general class of problems involving discrete interdependent risks with heterogeneous agents. There is a threat of an event that can only happen once, and the risk depends on actions taken by others. Any agent's incentive to invest in managing the risk depends on the actions of others. Security problems at airlines and in computer networks come into this category, as do problems of risk management in organizations facing the possibility of bankruptcy, and individuals' choices about whether to be vaccinated against an infectious disease. Surprisingly the framework also covers certain aspects of investment in R&D. Here we characterize Nash equilibria with heterogeneous agents and give conditions for tipping and cascading of equilibria.

Key Words: Nash equilibrium, tipping, cascading, terrorism, security, interdependence.

JEL Classification: C 72, D 80, H 23,

1 Introduction

Certain events can only occur once. Death is the obvious example: it is irreversible and unrepeatable. Extinction of a species takes this even further. More mundane examples are bankruptcy, being struck off a professional register for life, and other discrete events. There are in addition events that can in principle occur twice but that are so unlikely or so dreadful that one occurrence is all that can reasonably be considered. The events of 9/11/01 are of this type, as is a nuclear meltdown in a highly populated region. The probabilistic nature of events like these, together with the fact that the risk that one agent faces depends on the behavior of others, gives a unique and unnoticed structure to the incentives that agents face to manage these risks. For other recent papers dealing with the interdependence of security-related risks, see Keohane and Zeckhauser [26] and Orszag and Stiglitz [30]. Keohane and Zeckhauser show how individual responses to collective threats may undermine the effectiveness of government policies to combat them. Orszag and Stiglitz address the interdependency issue by showing that homeowners do not take into account the positive externalities associated with reducing damage to their neighbors when determining how fireproof a structure they should build. For reviews of the application of game theory to security and terrorism at the national level, see Sandler [32] and Sandler and Acre [33].

The key point for these problems is that an agent's incentive to invest in risk-reduction measures depends on how he expects others like him to behave. If he thinks that they will not invest in security, then this reduces the incentive for him to do so. But should he believe that they will invest in security, then it may be best for him to

do likewise. So there may be an equilibrium where no one invests in protection, even though all would be better off if they had incurred this cost, and indeed all incurring this cost may be an equilibrium. This situation does not have the structure of a prisoners' dilemma game, even though it has some similarities (see Kunreuther and Heal [27]). It contains elements of the coordination problem discussed by Heller [18], Crawford and Haller [7] and others, in that there may be many alternative equilibria of the system, some of which are Pareto ranked, but it is nevertheless not clear which equilibria will emerge. The interdependence between the strategies of agents means that in some cases there is a complementarity between them: by investing in risk reduction, one agent makes this strategy more attractive to others, so that the strategies work better when chosen by several agents than when chosen singly. In this sense some of our results appear to be similar to those in the literature on strategic complementarity (Bulow Geanakoplos and Klemperer [4]). However the discreteness of our strategy space and the fact that damages are non-additive - you only die once - means that the technical details are different.

It is possible that a change in the behavior of one agent can tip the system from one equilibrium to another (see Schelling [35]): a related phenomenon is cascading, when a change by one leads to a change by a second which provokes a change by a third, and so on (see Dixit [12] and also Farrell and Saloner [13]).

1.1 Features of the Problem

There are several different versions of this *interdependent security (IDS)* problem but all have certain features in common. We have already indicated one of these: a payoff

that is discrete. A bad event either occurs or does not, and that is the full range of possibilities. You die or you live. A firm is bankrupt or not. A plane crashes or it doesn't. You catch a disease or you do not. In these examples it is not useful to differentiate the outcomes more finely.

Another feature common to the problems that we consider is that the risk faced by one agent depends on the actions taken by others – there are externalities. The risk of an airline's plane being blown up by a bomb depends on the thoroughness with which other airlines inspect bags that they transfer to this plane. The risk that a corporate divisional manager faces that her company will be sent into bankruptcy depends not only on how she manages her divisional risks but also on how other division heads behave.

Finally there is a stochastic element in all of these situations. The question addressed is whether to invest in security when there is some probability, often a very small one, that there will be a catastrophic event that could be prevented or mitigated. This risk depends in part on the behavior of others, and the unfavorable outcome is discrete in that it either happens or does not.

These three factors – non-additivity of damages, dependence of risks on the actions of others, and stochasticity – are sufficient to ensure that there can be equilibria where there is underinvestment in risk-prevention measures. The precise degree of underinvestment depends on the nature of the problem. We focus initially on the two extremes that span the spectrum of possibilities. Both relate to security, one of airlines and the other of computer networks. If an airline accepts baggage that

contains a bomb, this need not damage one of its own planes: it may be transferred to another airline before it explodes. So in this framework one agent may transfer a risk fully to another. It may of course also receive a risk from another. There is a game of “pass the parcel” here. The music stops when the bomb explodes. It can only explode once, so only a single plane will be destroyed.

The structure of this game is quite different in the case of computer networks. Here it is commonly the case that if a virus (or hacker) enters the network through one weak point it (or he) then has relatively easy access to the rest of the network and can damage all other computers as well as the entry machine. Indeed many computer viruses are programmed to do this by sending themselves to all addresses in an infected machine’s address book. In this case the bad outcome has a characteristic similar to a public good: its consumption is non-rivalrous. Its capacity to damage is not exhausted after it has inflicted damage once. A bomb, in contrast, has a limited capacity to inflict damage, and this capacity is exhausted after one incident.

In both cases the incentives to take security measures depend on what others do. Suppose that there are a large number of agents in the system. In [27] we show that in the computer security problem, if none of the other machines are protected against viruses or hackers, then the incentive for any agent to invest in protection approaches zero as the number of agents increases. For airline security, if no other airline has invested in baggage checking systems and there is a high probability that bags will be transferred from one airline to another, in the limit the expected benefits to any airline from this investment approaches 63% of what it would have been in

the absence of contamination from others.

It is not only security problems that have this structure. It is common to all problems with discrete and interdependent risks. It applies to units of a multi-unit organization in which the risk of bankruptcy (a discrete event) faced by any unit is affected by its own choices and by the choices made by other units. In such a situation any unit's incentive to take actions to reduce bankruptcy risks is compromised by the knowledge that others are not being similarly diligent. A culture of risk-taking can spread through the organization because knowledge that a few groups are taking risks reduces the incentives that others have to manage them carefully.

Some decisions about research and development (R&D) investment also have this structure. The central issue here is that if several firms want to solve a problem, each may try on its own or may wait until another solves it first. The greater the probability that another will solve the problem first, the less the incentive to try to solve it oneself unless being first conveys an advantage such as a right to patent. With this type of interaction the externalities are negative rather than positive, in the sense that action by others makes action by oneself less attractive.

The problem of choosing whether or not to be vaccinated against an infectious disease has a similar structure. Firstly, vaccination is a yes-no choice and one can in general only catch the disease once. Secondly, the risk of catching the disease depends on the number of others who choose to be vaccinated, i.e. who invest in risk-management. So once again we have the structure of an IDS problem.

In the vaccination case, as in the R&D case, the externalities between agents are

negative in the sense that protection by others makes vaccination less attractive to oneself. In the security and bankruptcy models where externalities are positive, we have, as mentioned before, some of the properties associated with strategic complementarity: in the other cases we have something closer to strategic substitutability [4]. A good illustration of the complementarity case is provided by investment in visible burglar alarms: the decision by others in a neighborhood to protect themselves with alarms makes investment in protection more attractive to you because you are now more likely to be a target. "Weakest link" and "best shot" games for the provision of public goods (Hirschleifer [19], Sandler and Vicary [34], Vicary and Sandler [36]) also exhibit strategic complementarity or substitutability respectively and have points in common with our models: a generalization of our model to continuous strategy spaces¹ would have both types of game as special cases.

Our earlier paper [27] studied IDS problems where all agents are identical. Here we extend the analysis to the more general case of agents whose risks and costs differ and study the possibility of tipping. There may be one firm occupying such a strategic position that if it changes from not investing to investing in protection, then all others will find it in their interests to do the same. And even if there is no single firm that can exert such leverage, there may be a small group. We show when this can happen and how to characterize those agents having so much leverage that by switching policy they can change the equilibrium choices of all others, in the process introducing a measure of the leverage a firm can exercise over others. This

¹Currently under development.

is a measure of its strategic importance within the group.²

Obviously this finding has significant implications for policy-making. It suggests that there are some key players whom it is particularly important to persuade to manage risks carefully. Working with them may be a substitute for working with the population as a whole. They are in a certain sense leaders or trendsetters. We also show that equilibria in these models may be susceptible to cascading, in that a change by a first firm can lead to emulation by a second, and the actions of these two can lead to emulation by a third, and so on.

In our earlier paper the probabilities describing the risks of loss were taken as exogenous. That was a simplification. It does not apply to many deliberate acts such as terrorism. Take the case of airline security. In practice terrorists will try to attack the airlines with the weakest security records. So if one airline improves its security then this will reduce the chances of an attack on it and increase the chances of attacks on others. Probabilities often respond to the policies adopted by the agents (Sandler [32], Woo [38]). We model this phenomenon here. The tragic fate of PanAmerican's flight 103 in 1988 illustrates this and other points from this introductory discussion. The flight was destroyed by a bomb loaded onto Malta Airlines at Gozo, Malta, flown to Frankfurt and then transferred to PanAm in London. Malta Airlines and Gozo were presumably chosen because they were seen as having weak security procedures relative to PanAm in London, and in the knowledge that for cost and logistical reasons

²Tipping can also occur with identical agents, the case considered in our earlier paper [27], but in that case if one agent can tip an equilibrium then so can any. All agents have the same leverage over others, the same strategic importance.

inter-airline baggage is never screened.³

The next two sections of the paper develop an IDS model where the probabilities and risks differ between agents and then characterizes the structure of the Nash equilibria. Section 4 then considers tipping and cascading, introducing a measure of a firm's leverage over others and showing that in some cases those with the greatest leverage are those which produce the greatest aggregate negative externalities: if you can convince them to invest in security other agents are likely to follow suit. After considering the case of endogenous probabilities in Section 5, we turn in Sections 6-9 to a set of other IDS problems to see how they differ in structure from the airline security problem. We begin with computer security then turn to bankruptcy, investment in R&D and finally to vaccinations. The concluding section summarizes the findings and suggests directions for future research. An appendix contains formal proofs of the results and also proves the existence of a Nash equilibrium in pure strategies for the models considered here. Because of the discreteness of the strategy space, the standard proof of existence as given by Nash in his classic article does not apply.

2 The Model

Initially we think in terms of the security of airlines, as this is an example that is both topical and canonical. There are $n \geq 2$ separate airlines. During the course of a given time period the airline makes a certain number of trips, each of which is

³We believe that El Al is the only airline to screen bags received from other airlines.

identical. Consider a given plane trip initiated by airline i . Assume that the airline has made no investments in security systems. Let p_{ij} be the probability that on any trip a bag containing a bomb is loaded onto airline i and is then transferred to airline j and explodes on j . If $i = j$, we have the probability that an airline loads a bag with a bomb and this explodes on its own plane. We denote by $p_i = \sum_j p_{ij}$ and by $\tilde{p}_i = \sum_{j \neq i} p_{ij}$. Thus p_i is the probability of airline i loading a bomb that explodes and \tilde{p}_i is the probability that it loads a bomb that explodes on another airline- a measure of the risk that it poses to others. We expect that $p_i < 1$ so that there is some chance that the airline does not load a bag with a bomb that explodes. Each airline can either invest in a security system S at a cost per trip of $c_i > 0$ or not invest N . Security systems are assumed to be completely effective so that they eliminate the chance of a bomb coming through the airline's own facility. In the event that a bomb explodes on a plane the loss is $L > 0$. The initial income of an airline is $Y > c_i \forall i$.

In the case of just airlines A_1 and A_2 maximizing expected profits this framework gives rise to the following payoff matrix showing the outcomes for the four possible combinations of N and S . If both airlines invest in security systems then their payoffs per trip are just their initial incomes net of the investment costs. If A_1 invests and A_2 does not, then A_1 has a payoff of income Y minus investment cost c_1 minus the expected loss from a bomb transferred from A_2 that explodes on A_1 (i.e, $p_{21}L$), while A_2 has a payoff of income Y minus the expected loss from a bomb loaded and exploding on to its plane, $p_{22}L$. If neither invests then A_1 has a payoff of income Y minus the expected loss from a bomb loaded and exploding on to its own plane

$p_{11}L$ minus the expected loss from a bomb transferred from A_2 , that explodes on A_1 (i.e., $p_{21}L$) conditioned on there being no explosion from a bomb loaded by A_1 itself ($1 - p_{11}$). A_2 's payoff is determined in a similar fashion.

A_1/A_2	S	N
S	$Y - c_1, Y - c_2$	$Y - c_1 - p_{21}L, Y - p_{22}L$
N	$Y - p_{11}L, Y - c_2 - p_{12}L$	$Y - p_{11}L - (1 - p_{11})p_{21}L, Y - p_{22}L - (1 - p_{22})p_{12}L$

Choosing to invest in security measures is a dominant strategy for 1 if and only if

$$c_1 < p_{11}L \text{ and } c_1 < p_{11} [1 - p_{21}] L \quad (1)$$

The condition that $c_1 < p_{11}L$ is clearly what we would expect from a single airline operating on its own. The tighter condition that $c_1 < p_{11} [1 - p_{21}] L$ reflects the risk imposed by a firm without security on its competitor: this is the risk that dangerous baggage will be transferred from an unsecured airline to the other. This negative externality plays a critical role in our analysis and we need to understand its structure as the analysis is expanded to cover n airlines.

Let $X_i(n, K)$ be the expected negative externality from all other airlines to airline i when airlines in the set K invest in security and there are n airlines in total. For

four firms and $i = 1$ we have terms of the form

$$X_1(4, \{2, 3, 4\}) = 0$$

$$X_1(4, \{3, 4\}) = Lp_{21}$$

$$X_1(4, \{2, 4\}) = Lp_{31}$$

$$X_1(4, \{4\}) = L\{p_{21} + (1 - p_{21})p_{31}\}$$

$$X_1(4, \{3\}) = L\{p_{21} + (1 - p_{21})p_{41}\}$$

$$X_1(4, \{2\}) = L\{p_{31} + (1 - p_{31})p_{41}\}$$

$$X_1(4, \emptyset) = L\{p_{21} + (1 - p_{21})p_{31} + (1 - p_{21})(1 - p_{31})p_{41}\}$$

In the Appendix we show that this last expression, for $X_1(4, \emptyset)$, can readily be derived from the event tree corresponding to the four agent problem. In all cases when transfers from more than one airline are possible then the losses from transfers from the second and subsequent firms have to be conditional on there being no losses from previous transfers. For n airlines when none of them invest in security this generalizes to the following formula for the externality inflicted on the first:

$$X_1(n, \emptyset) = L \sum_{j=2}^{j=n} p_{j1} \prod_{k=2}^{k=j-1} (1 - p_{k1})$$

where it is understood that $\prod_{k=2}^{k=j-1} (1 - p_{k1}) = 1$ when $j = 2$. If firms in the set K are investing in security then the total externality to firm 1 is given by an extension of

this formula, replacing \emptyset by K and noting that $p_{kj} = 0 \forall k \in K$:

$$X_1(n, K) = L \sum_{j=2}^{j=n} p_{j1} \prod_{k=2}^{k=j-1} (1 - p_{k1}) \quad (2)$$

The condition for S to be a dominant strategy for firm 1 when there are n firms with none investing is that

$$c_1 < p_{11} [L - X_1(n, \emptyset)] = c_1(n, \emptyset) \quad (3)$$

Here $c_1(n, \emptyset)$ is the maximum cost to agent 1 consistent with S being the best strategy for 1 when no other firms invest in security. More generally

Definition 1 $c_i(n, K)$ is the maximum cost of investment in security at which agent i will choose to invest in security when there are n agents and those in the set K have already invested in security.

Clearly $X_1(n, \emptyset) > X_1(n, \{2\}) > X_1(n, \{2, 3\}) > \dots > X_1(n, \{2, 3, 4, \dots, n-1\})$ so that $c_1(n, \emptyset) < c_1(n, \{2\}) < \dots < c_1(n, \{2, 3, 4, \dots, n-1\})$. This implies that as we add more agents who do not invest in security the externality on any other agent increases and the condition for them to want to invest in security becomes more demanding and so such investment becomes less likely. We showed in [27] that if all agents are identical then $\lim_{n \rightarrow \infty} X_1(n, \emptyset) = L(1 - e^{-p})$ where p is the common probability of a transfer. From now on, we will drop the argument n from expressions for X and c , as in general n will be held fixed throughout the analysis.

3 Nash Equilibria

The nature of the Nash equilibrium in the interdependent security model naturally depends on the parameters. From the payoff matrix it is clear that (S, S) is a Nash equilibrium if $c_i < p_{ii}L$ and is a dominant strategy if $c_i < p_{ii}L(1 - p_{ji})$ where i and j are 1 or 2. (N, N) is a Nash equilibrium if $c_i > p_{ii}L(1 - p_{ji})$ and a dominant strategy if $c_i > p_{ii}L$. From these inequalities we note that (S, S) and (N, N) are both Nash equilibria if $p_{ii}L(1 - p_{ji}) < c_i < p_{ii}L$. Finally if $c_1 > p_{11}L$ but $c_2 < p_{22}L(1 - p_{12})$ then (N, S) is a Nash equilibrium, and if 1 and 2 are interchanged then the equilibrium is (S, N) . This configuration of Nash equilibria is summarized in Figure 1. Note that if $c_1 = c_2$ then we are on the diagonal of figure 1 and the only possible equilibria are (S, S) , either (S, S) or (N, N) , and (N, N) . In this case mixed equilibria are not possible, as stated in our earlier paper [27].

Figure 1 shows that even for two agents there is a wide variety of Nash equilibria for this problem including cases where there are two possible equilibria $\{S, S\}$ and $\{N, N\}$. As one expands the number of agents the number of possible equilibria expands exponentially.

Recall that $c_i(\emptyset)$ is the maximum cost at which agent i will invest in security if no others are investing. Clearly if $c_i > c_i(\emptyset) \forall i$ then $\{N, N, \dots, N\}$ is a Nash equilibrium. More generally we can characterize a Nash equilibrium as follows:

Definition 2 *A Nash equilibrium is a possibly empty set E of agents choosing S such that $c_i < c_i(E) \forall i \in E$ and $c_i > c_i(E) \forall i \notin E$.*

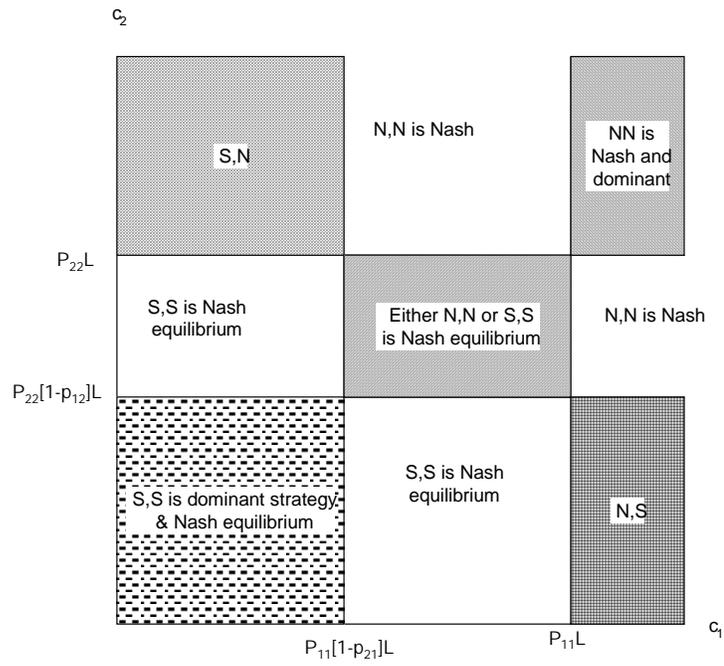


Figure 1: Nash equilibria as a function of c_1 and c_2 .

In words, a Nash equilibrium is a situation where some firms are choosing S and some N , and for those choosing S the actual cost of investing in security is less than the maximum that is justifiable economically, given the choices of others, and for those choosing N , the actual cost is greater.

In the appendix we prove that a Nash equilibrium in pure strategies exists for the general model of this section.

4 Tipping & Cascading

In some cases a change of strategy by one agent or a small set of agents can shift the equilibrium radically. We refer to this change as *tipping* in the sense of Schelling [35], Katz and Shapiro [22], Watts [37] (in the context of general networks) and more recently Gladwell [14]. For example, there may be a Nash equilibrium at which no agent invests in security. Yet if one agent changes strategy and invests - possibly in response to events or incentives outside the game - then all other agents may follow suit. We illustrate how tipping can occur for the case where all agents initially choose strategy $\{N, N, \dots, N\}$: we give conditions for a change by a subset of these agents to lead all the others to follow suit and invest in security, producing an equilibrium $\{S, S, \dots, S\}$. Of course an equilibrium cannot be tipped from N to S if it is an equilibrium in dominant strategies, so that the dominant strategy equilibria in figure 1 could not be tipped. These include the (N, S) and (S, N) equilibria. If we have an equilibrium at which N is a dominant strategy for some firms and not for others then this could be tipped provided that the firms for which N is dominant are included in

those whose strategies are exogenously altered as part of the tipping process, e.g. by being taxed or some other policy change.

Intuitively it seems that there are two important aspects of the tipping phenomenon. One is the vulnerability of an agent to being tipped from not investing to investing, which depends on how close its cost is to the maximum cost at which investment is justified. If this gap is small then a small change in the externalities imposed on the agent by others may suffice to change its choice of strategy. The second important aspect of tipping is the change in the externalities imposed on other agents when one agent changes its policy. An agent for which this change is big is more likely to cause tipping than one for which this is small. In general the possibility of tipping depends on both of these factors - on how close agents are to changing their strategy choices and how large the negative externalities are from some agents on others because they do not invest in protection.

Consider a Nash equilibrium where all agents choose strategy N . A *critical coalition* K is a group of firms that by switching from N to S can tip the equilibrium to one where all firms invest. It is a *minimal critical coalition (MCC)* if it is a critical coalition and no subset is a critical coalition.⁴ Formally consider a Nash equilibrium such that $c_i > c_i(\emptyset) \forall i$, so no agents invest in security, and $c_i < c_i(K) \forall i \notin K$. The firms in K form a critical coalition: if they switch from N to S , then this leads all other firms to follow suit because their critical costs will now be greater than their actual costs of investment.

⁴The concept of minimum critical coalition is introduced in Heal [16]. The minimum critical coalition is in general not unique, although it is in the particular case considered in proposition 2.

Firms in an MCC are an important group. If they change from not investing to investing, then all others follow suit. The reason this occurs is that when these agents invest they reduce the externalities on others sufficiently that it is now cost-effective for the others to invest in security too. To understand the impact that an agent has on others we need to know the total externalities that each agent generates on all other agents by not investing in security.

We define the externality from i to j when the firms in set K , $i \notin K$, are investing in security as the change in the total externality to j when i switches from not investing to investing. Denote this by $\Delta_{ij}(K)$. Formally this is

$$\Delta_{ij}(K) = X_j(K) - X_j(K + i), i \notin K$$

From the definitions of $X_j(K)$ and $X_j(K + i)$ we have:

$$\Delta_{ij}(K) = Lp_{ij} \prod_{k=2}^{k=j-1} (1 - p_{kj}) \quad (4)$$

where as before $p_{ij} = 0 \forall i \in K$. The total externality generated by i is just the sum over j of (4):

$$\Delta_i\{K\} = \sum_j \Delta_{ij}(K) = L \sum_{j \neq i} p_{ij} \prod_{k=2}^{k=j-1} (1 - p_{kj}) \quad (5)$$

This is a reasonably compact and intuitive expression for the total externality generated by i when the firms K are investing in security. It is the loss from a single occurrence times the sum of the probabilities of a transfer from i conditioned on the other firms outside K not having inflicted damage on a firm already. We focus

on the case when no firms are investing in security, in which case the appropriate index is $\Delta_i \{\emptyset\}$. With these definitions in place, we can now give the following formal characterization of a critical coalition.

Proposition 1 *A critical coalition is a set of agents K such that*

$$p_{jj} \sum_{k \in K} \Delta_{kj} = c_j - p_j (L + X_j(\emptyset)) \forall j \notin K$$

The left hand side here is the reduction in the total externality imposed on agent j when agents in K switch from N to S , multiplied by p_{jj} . It can be viewed as the expected benefit to j of having agents in the critical coalition invest in security. The right hand side is derived from equation (3) and is the difference between the actual cost of investment in security and the maximum that it is worth paying to invest. The derivation of this inequality is almost immediate from (3) and the definition of Δ_{kj} . It juxtaposes the two issues referred to above - the impact of a change in policy by one agent on the externalities faced by others, and the nearness of these others to changing their strategy choices.

Some firms - those for which $\Delta_i \{\emptyset\}$ is large - are clearly more likely, in some general sense, to cause tipping than others.

In general there is no easy way of characterizing the agents who have greatest leverage. There is however one interesting case in which this is possible, which is when the contagion probabilities p_{ij} are the same for all j for a given i . This is the case in which an agent is equally likely to transfer a bag to any other agent,

so that $p_{ij} = p_{ik} \forall j, k \neq i$. This case is implausible for airlines, but in some of the applications considered later is quite realistic - for example the case of bankruptcy in a multidivisional firm, or the case of a computer network. In the terminology of Kearns and Mansour [24] this assumption makes the IDS game a summarization game. Kearns and Ortiz [25] show that in this case the Nash equilibria of an IDS game can be computed in a time that is a polynomial function of the number of parameters, whereas in the general case the game is NP-complete. When $p_{ij} = p_{ik} \forall j, k \neq i$ the agents with greatest leverage are those with the greatest values of $\Delta_i \{\emptyset\}$, that is, those that create the greatest externalities for others.⁵ These are the firms which will “tip” a Nash equilibrium from not investing to investing.

Proposition 2 *Assume that $p_{ij} = p_{ik} \forall j, k \neq i$. If there is a minimal critical coalition of $k < n$ agents then it must consist of the first k agents ranked by $\Delta_i \{\emptyset\}$.⁶*

Proof. The proof is in the Appendix. ■

Note that in this case a minimum critical coalition is unique. Proposition 1 shows that if there is a minimal critical coalition, then it consists of agents who impose the largest externalities on the others. There is a simple intuition for this result. The decision to invest in protection is determined by the expected direct reduction in damage (i.e. $p_{ii}L$) minus the likelihood that the agent will be harmed by unprotected agents multiplied by the resulting loss L . Order the agents so that

⁵In general the distribution of the externalities matters as well as the total in calculating leverage. However when all agents have the same chance of being impacted negatively by others only the total matters.

⁶*Ties will be broken randomly if needed. The ranking of agents by $E_i \{\emptyset\}$ is of course a function only of the parameters of the model and is not affected by the strategy choices of agents.*

agent 1 has $\max\{\Delta_i\{\emptyset\}\}$, agent 2 has the second highest value, and so on. Agent 1 inflicts the highest expected harm on others. Hence by inducing it to invest in protection one has the most impact on the incentive of other unprotected agents to invest in protection. Should a change of strategy by agent 1 not be sufficient to do this, then one has to convince agent 2 to invest in protection as well in the hopes that this will lead the remaining unprotected agents to invest. The smallest number of agents to induce this type of tipping behavior is deemed an MCC.

4.1 Numerical Examples of Tipping and Cascading

The following numerical example demonstrates that such an MCC can indeed exist. We shall consider three airlines and let 1 and 2 be identical. The characteristics of these airlines are such that the only Nash equilibrium is one where none of them invest in security. Yet if airline 3 changes from not investing to investing - perhaps as a result of a financial incentive or regulatory pressure or some other factor outside of the model - then both others will change as well and there is a new equilibrium at which all are investing. The change was produced by the change in 3's behavior.

Let $p_{1j} = p_{2j} = 0.1$ and $p_{3j} = 0.5$ and $L = 1000$. In addition $p_{11} = p_{22} = 0.1$ and $c_1 = c_2 = 85$. We do not specify p_{3j} or c_3 . In this setting

$$c_1(\emptyset) = 0.1 \left[1000 - \frac{(0.5)(1000)}{2} - \left(0.1 - \frac{0.05}{2} \right) \frac{1000}{2} \right] = 71.25$$

As $c_1 = c_2 = 85 > c_1(\emptyset) = c_2(\emptyset) = 71.25$, neither firm 1 nor firm 2 will invest in security if no other firm is investing. And we can clearly choose c_3 so that it is large

enough that firm 3 will not invest either and (N, N, N) is the Nash equilibrium. And if firm 3 does not invest, then not investing is a dominant strategy for both the other firms for any cost above 75.

Suppose that for some reason airline 3 changes policy and invests. It now imposes no externality on the other firms and so does not affect their decisions. To understand the choices of firms 1 and 2 we simply have to apply inequality (1), which gives a critical cost level of 90, meaning that investment will now be a dominant strategy when the cost is less than 90. As the actual cost for firms 1 and 2 is less than this by assumption at 85, we see that after firm 3 has changed strategy from N to S for both firms 1 and 2 the dominant strategy has changed from not investing to investing. Airline 3 therefore has the capacity to *tip* the equilibrium from not investing to investing by changing its policy. It is easy to verify that airline 3 imposes the largest externalities on the other airlines in accordance with Proposition 2 above.

The tipping phenomenon is shown geometrically in the following two diagrams. These are similar to figure 1 above, showing the sets of $\{c_1, c_2\}$ values corresponding to different equilibrium types. The key point in seeing tipping geometrically is that this diagram for firms 1 and 2 depends on what firm 3 does. A change by 3 alters the entire equilibrium diagram for the other two firms.⁷ When firm 3 does not invest, as in figure 2, not investing is a dominant strategy for the other firms as their cost point $(85, 85)$ lies in the quadrant bounded below by $(75, 75)$. When firm 3 changes and invests, then the whole diagram for the other firms alters, now looking as in figure 3.

⁷We are really looking at a three-dimensional version of figure 1, and the diagrams for firms 1 and 2 are slices through this for different strategy choices for firm 3.

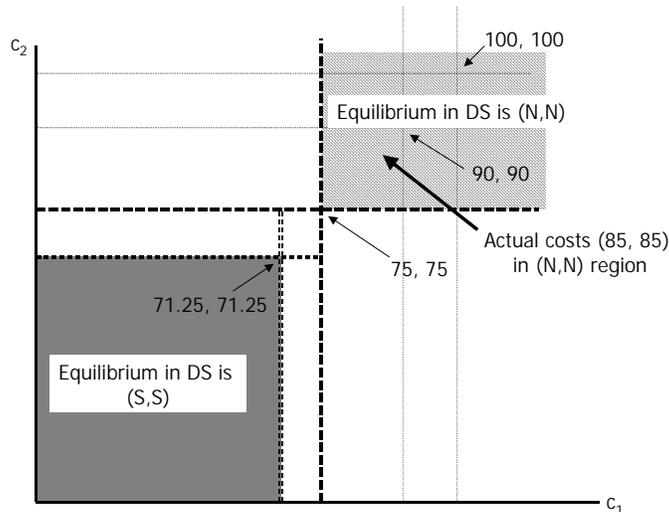


Figure 2: Equilibria for firms 1 and 2 when 3 does not invest and imposes externalities on them. In this case $(85, 85)$ is in the region in which not investing is a dominant strategy.

The region in which investing is a dominant strategy is now greatly enlarged because of the removal of the externalities generated by 3 and includes the point $(85, 85)$ so that it includes the point representing firms 1 and 2.

The tipping phenomenon that we are characterizing here is in fact more general than the particular illustrative context as indicated by the following question: Given a Pareto inefficient Nash equilibrium in a general game, does there exist a subset of agents who by changing their strategy choices can induce all others to alter their strategy choices in such a way that the new outcome is efficient? The previous proposition and example show that this is the case in the interdependent security problem. It would be interesting to ask this question for a broader class of games. Kearns and Ortiz [25] use computational algorithms to give further numerical examples of minimum critical coalitions and tipping in the context of data from 49 major international

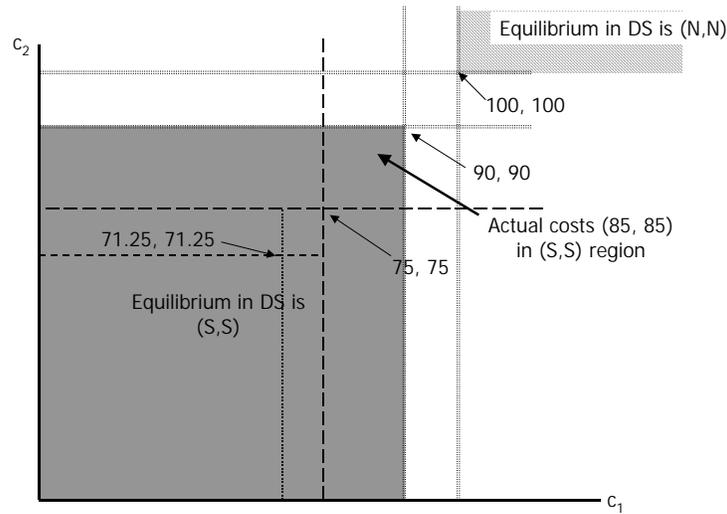


Figure 3: Equilibria for firms 1 and 2 when 3 invests and imposes no externalities on them. In this case $(85, 85)$ is in the region in which investing is a dominant strategy.

airlines.

Our model can also give rise to the phenomenon of cascading (see also Dixit [12]). This refers to a situation where one firm changes its policy, and this leads another to follow suit. The fact that two firms have changed now persuades a third to follow, and when the third changes policy this creates the preconditions for a fourth to do so, and so on. The analogy with a row of dominoes is compelling: the first knocks down the second, which knocks down the third, and so on. To see how this can happen in the our model, suppose that we have a Nash equilibrium at which all airlines choose N and assume in addition we can number firms $1, 2, 3, \dots$ so that the following conditions are satisfied:

- When 1 switches from N to S then 2's best strategy changes from N to S but no other firm's best strategy changes

- When 1 and 2 have switched from N to S then 3's best strategy changes from N to S and no other firm's best strategy changes.
- When 1, 2 and 3 have switched from N to S then 4's best strategy changes from N to S and no other firm's best strategy changes.

or in general

- When 1, 2, 3, ..., J have switched from N to S then $(J + 1)$'s best strategy changes from N to S and no other firm's best strategy changes for all firms $J > 1$.

If such an ordering of the firms exists then it is immediate that if firm 1 switches from N to S then it will start a cascade in which 2 changes followed by 3 then by 4 etc. etc. We can readily modify the numerical example above to illustrate this cascading process. Specifically, keep the probabilities as above and let $c_1 = 85$ as before but $c_2 = 95$. Then it is immediate from figure 2 and 3 that (c_1, c_2) is in the region where (N, N) are the dominant strategies when three does not invest but also is in the region where (N, S) is the equilibrium when three does invest (see also figure 1). So in this case when three changes from N to S this causes two to change from not investing to investing as well. But once firms two and three are investing, firm one is effectively on its own and will invest if $c_1 < p_{11}L = 100$, which is satisfied. So when two follows three and changes from not investing to investing it will cause one to follow suit, generating a cascade.

5 Endogenous Probabilities

So far the risks faced by the airlines are assumed to be independent of their behavior. In reality if some airlines are known to be more security-conscious than others, they are presumably less likely to be terrorist targets. There is a resemblance here to the problem of theft protection: if a house announces that it has installed an alarm, then burglars are likely to turn to other houses as targets [27]. In the case of airline security, terrorists are more likely to focus on targets which are less well protected, so that the p s depend on the investment in security. This is the phenomenon of displacement or substitution, documented in Sandler [32]. Keohane and Zeckhauser [26] also consider the case of endogenous terrorist risks.

We assume here that the risk faced by an airline not investing in security increases as the fraction of airlines investing in security increases. In other words, if more airlines from a given population invest in security then those who do not invest become more vulnerable. Formally let $\#K$ be the number of airlines investing in security, i.e. the number in the set K . The relevant probabilities facing those firms not investing in security, $p_{ij}(\#K)$, are increasing in $\#K$. For airlines that have invested in security the p s are assumed to be independent of $\#K$.

Now return to equation (3) above, defining the cost of investment that marks the boundary between a firm i investing and not investing in security when no other firm invests when p_{ij} are exogenous:

$$c_i < p_{ii} [L - X_i(\emptyset)] = c_i(\emptyset)$$

As an increasing number of other firms invest in security, then for a non-investing firm the probability p_{ij} will increase. Hence the value of $X_i(K)$ will change. The expression for the critical value of the cost of investment for airline i is thus :

$$c_i(K) = p_{ii}(\#K) L \left[1 - \sum_{j \notin K} \frac{p_{ij}(\#K)}{n-1} \prod_{k < j, k \notin K} \left(1 - \frac{p_{ik}(\#K)}{n-1} \right) \right] \quad (6)$$

The right hand side of (6) increases in $\#K$ via p_{ii} but also depends on $\#K$ through the p_{ik} s that enter into the expression for the externality imposed on i . The sign of the impact of a change in $\#K$ on the externality is not clear *a priori*: an increase in the number of firms investing will raise $\sum_{j \notin K} \frac{p_{ij}(\#K)}{n-1}$ but will also decrease

$$\prod_{k < j, k \notin K} \left(1 - \frac{p_{ik}(\#K)}{n-1} \right).$$

We assume, as seems generally reasonable, that the total externality imposed on any non-investing firm decreases as the number of investing firms increases, in which case an increase in $\#K$, the number of firms investing, will increase the right hand side of (6) and raise the value of $c_i(K)$. This means that an agent is more likely to invest in security for the case where probabilities are endogenous than when these probabilities are exogenous. This assumption also implies the analysis with constant probabilities remains qualitatively valid in the endogenous case.

Of course the computation of $c_i(K)$ is now more complex due to the dependence of the probabilities on the number of firms investing in security. The endogeneity of probabilities should lead more firms to invest in protection given that they are now more likely to be targets. The concept of a minimum critical coalition also carries over unaltered to the world of endogenous probabilities, although the actual MCCs

will very likely be different.

Given that the basic concepts do not change qualitatively, Propositions 1 and 2 on tipping are relevant to a model with endogenous probabilities. It should now be easier for a coalition to tip the other firms into investing for the following reason: not only does a decision by a firm to invest reduce the externalities but it also increase the risk that a firm who did not invest in security will become a target. In the next four sections we examine how the IDS model applies to a set of problem contexts where damage are non-additive and there are negative stochastic externalities. To keep the analysis simple the probabilities are assumed to be exogenous.

6 Computer Security

When a virus affects a computer (the equivalent of a bag with a bomb being loaded by an airline) it can be transmitted to all other computers on the network and can damage them all rather than just one of them (Anderson [2]). Let p_i be the probability that computer i is infected by a virus and \tilde{p}_i be the probability that it is infected by a virus and this is transmitted to *all* other computers. Clearly $\tilde{p}_i \leq p_i$ and the ps do not refer to independent events. The other notation is the same as in the airline security problem. The stochastic negative externalities for the case of four computers

are given by terms that include the following:

$$X_1(4, \{2, 3, 4\}) = 0$$

$$X_1(4, \{3, 4\}) = L\tilde{p}_2$$

$$X_1(4, \{4\}) = L\{\tilde{p}_2 + (1 - \tilde{p}_2)\tilde{p}_3\}$$

$$X_1(4, \emptyset) = L\{\tilde{p}_2 + (1 - \tilde{p}_2)\tilde{p}_3 + (1 - \tilde{p}_2)(1 - \tilde{p}_3)\tilde{p}_4\}$$

and in the general case of n computers when none of them invest in security this generalizes to the following formula for the externality inflicted on the first:

$$X_1(\emptyset) = L \sum_{j=2}^{j=n} \tilde{p}_j \prod_{k=2}^{k=j-1} (1 - \tilde{p}_k)$$

where it is understood that $\prod_{k=2}^{k=j-1} (1 - \tilde{p}_k) = 1$ when $j = 2$.

If agents in the set K are investing in security then the total externality to agent 1 is given by

$$X_1(K) = L \sum_{j \notin K} \tilde{p}_j \prod_{k < j, k \notin K} (1 - \tilde{p}_k) \quad (8)$$

The condition for S to be a dominant strategy when there are n agents with none investing is that

$$c_1 < p_1 [L - X_1(\emptyset)] \equiv c_1(\emptyset) \quad (9)$$

Here as before $c_1(\emptyset)$ is the maximum cost to agent 1 consistent with S being a Nash equilibrium when no other agents invest in security. In K-H [27] we show that if all agents are identical the term $X_1(\emptyset)$ goes to L as $n \rightarrow \infty$. The proof used there can

be modified to apply to the present case, so that it is again the case that $X_1(\emptyset)$ goes to L as $n \rightarrow \infty$.

The definitions of Nash equilibrium and minimum critical coalition carry over unchanged from the previous sections. Now it is natural to assume, as we have, that the contagion probabilities are uniform, so that agents' leverage can be calculated by the value of the externalities that they impose on others when they switch policy. We can therefore prove an exact analog of proposition 2 for the computer network case:

Proposition 3 *A minimal critical coalition of $k < n$ agents must consist of the first k agents ranked by $E_i\{\emptyset\}$ or equivalently by \tilde{p}_i .*

The proof is exactly as before, and we can use the same numerical example to illustrate the proposition. In this case we find that $c_1(\emptyset) = 45$ and $c_1(\{3\}) = 90$ so that if computer 1 switches policy, it tips the network from not investing to investing in security.

7 Bankruptcy of Firms

Consider a multi-divisional organization, such as an investment bank, in which each division has some degree of decision-making autonomy and can incur risks on behalf of the entire organization. If any one division miscalculates grossly, incurring a large risk that causes a catastrophic loss, it may force the entire organization into bankruptcy. Several years ago the British merchant bank Barings, at that point the longest-established bank in the UK, was destroyed by the actions of a single trader in its Singapore branch. Nick Leeson incurred positions that put at risk sums that could

and indeed did destroy the company.⁸ In a rather different line of business Arthur Anderson was recently sent into bankruptcy in large part by the actions of its Houston branch in managing the Enron audits. Union Carbide suffered catastrophic losses from the accident at Bhopal in 1984 that eventually led to the firm being bought by Dow Chemical.

In each of these cases the situation is analytically similar to the computer security problem. An organization consists of a group of divisions $i = 1, \dots, n$, each of which can incur risks for which the company as a whole is liable. Let p_i be the probability that division i incurs a loss so that management closes down only this division and \tilde{p}_i be the probability that division i incurs such a large loss that the entire company is bankrupt and every division is closed. As in the computer security case $\tilde{p}_i \leq p_i$ and the p_i do not refer to independent events. The loss to a division in the event of its being closed is L . One should view L as the costs that employees of the division will incur if their division or the entire firm goes bankrupt. These include the search costs for new employment and other negative features associated with losing ones job including loss of reputation. Divisions can invest in monitoring their risks at a cost c_i , so they can avoid the loss L .⁹

Clearly when division i takes on a risk, it is imposing an external effect on other divisions because there is some chance that a large loss to this division will cause the firm to be closed down. Nick Leeson in Barings imposed risks on all branches of

⁸For a more detailed description of the factors causing the collapse of Barings Bank see chapter 1 of Hoch and Kunreuther [20].

⁹One could also think of c_i as an opportunity cost resulting from avoiding certain deals that might compromise the entire enterprise.

Barings, and Anderson's Houston branch similarly imposed risks on all of Anderson. And as before these losses are non-additive: the risk is only relevant if the other divisions have not already been closed down by losses originating elsewhere. So the problem is identical in structure to the computer network problem. The total external costs imposed on division 1 when no other divisions are managing risks are given by

$$X_1(\emptyset) = L \sum_{j=2}^{j=n} \tilde{p}_j \prod_{k=2}^{k=j-1} (1 - \tilde{p}_k)$$

In the present context this means that the incentive that any division faces to invest in risk-control depends on whether others are making similar investments. Senior management may want each division to invest in loss prevention. Due to the negative externalities, divisions may be loathe to incur these costs because of their adverse impacts on divisional profits. From the perspective in overseeing the entire firm, senior management will seek policy measures that will change the payoffs and make investing in risk-control a dominant choice, but it may be difficult for them to do this if each of the divisions operates in a decentralized manner. There will also be the possibility of tipping the equilibrium by persuading a small number of divisions to adopt stricter controls, with an exact analog to proposition 2. For a more extensive analysis of the managerial and organizational implications of this problem, see [28].

8 Investing in Research and Development

The same IDS structure has relevance to the problem of determining whether to invest funds in research and development (R&D), a topic on which there is an extensive

literature (see e.g. Dasgupta and Stiglitz [8] and [9], Dixit [11] and Grossman and Shapiro [15]). In this literature the concern is to characterize the privately and socially optimal levels of investment in R&D and the relationship between them, which typically depends on institutional structures such as patent rights. Here we explore the investment levels that are privately optimal, and show that under the assumption that the investment in R&D is discrete - a firm invests or does not invest - the problem has the IDS structure, which enables us to give a more complete characterization of the Nash equilibria than the earlier models. We also explore how the strength of intellectual property rights affects the equilibrium.

Consider a group of firms, each of whom are trying to solve the same problem or trying to discover the same new facts. If one firm solves the problem or discovers the facts, then its solution may be available to some or all others at no cost or at a very low cost. In such a situation each firm has to decide whether to invest in obtaining the information or making the discovery, bearing in mind that another firm might make the discovery. If the information from the other firm were freely available, any investment of its own would be redundant.

The investment decision here has the same formal structure as the problems considered in previous sections, with one important difference. In this case investment decisions are not mutually reinforcing. In the airline, computer security and bankruptcy cases, investment by one agent increases the incentive for others to invest and it is this property that can lead to tipping behavior or cascading behavior. A move by one firm may encourage other firms to do the same and could start an avalanche. In

the R&D case, investment by one firm *discourages* others from following suit. Each firm knows that as the number of firms attacking a particular problem increases, the chances that the problem will be solved by one of the others increases, so making it more efficient to wait until one of these others discovers the solution. The R&D problem has the same formal mathematical structure as the other IDS problems discussed above, except for a difference in the sign reflecting the interactions between firms. This means that we no longer expect to see the tipping or cascading of equilibria.

Assume that firm i can invest in R&D at a cost of c_i . This generates a payoff of G with probability p_i . There is, in addition, a chance p_j that another firm j invests and succeeds, in which case the information it gains reaches firm i . If I stands for investing and N for not investing then the payoff for the two by two case is

Payoff matrix for firms 1 and 2 in the R&D problem

1/2	I	N
I	$Y - c_1 + p_1 G + (1 - p_1) p_2 G, Y - c_2 + p_2 G + (1 - p_2) p_1 G$	$Y - c_1 + p_1 G, Y + p_1 G$
N	$Y + p_2 G, Y - c_2 + p_2 G$	Y, Y

Here if neither invests then there is no chance of either getting the information and so both their payoffs are their initial income Y . If firm one invests and two does not, then the payoff to the investor is $Y - c_1 + p_1 G$, income net of the cost of investing plus the expected gain from the investment. The payoff to the non-investor here is $Y + p_1 G$, income plus the expected gain as the information is transferred to it from the successful investor. Finally if both invest then firm i has a payoff of $Y - c_i + p_i G + (1 - p_i) p_j G$, which is income net of the cost of investment plus the

expected gain from its own investment plus the expected gain from the other's investment conditional on its own investment not having succeeded.

In this payoff matrix I is a dominant strategy if and only if

$$c_i < p_i [1 - p_j] G \tag{10}$$

Thus the possibility of getting the information free from someone else reduces the incentive to invest in R&D: without this possibility the equivalent inequality would obviously be $c_i < p_i G$. The term $[1 - p_j]$ represents what was previously labeled contagion in [27]. In this context it might be called the *free rider* effect since there is a temptation for each firm to take advantage of the other firm's R&D investment. The knowledge that firm j is investing will reduce the incentive that firm i has to do likewise.

8.1 Nash Equilibrium

The Nash equilibrium for this problem differs from the airline and computer security cases because there is less incentive to invest in R&D if others have already done so. If no firms are investing then the return from investment is at its highest level while if all other firms are investing then the expected returns from investment is at its lowest level. We work with the simplest possible case of two firms where there is no advantage of being the first to discover: the results generalize readily. We initially suppose the firms to be different and show how this simplifies when firms are identical.

We know already from (10) that (I, I) is a Nash equilibrium if

$$c_1 < p_1 [1 - p_2] G \text{ and } c_2 < p_2 [1 - p_1] G$$

Similarly (I, N) is a Nash equilibrium if

$$p_1 G > c_1 \text{ and } c_2 > G p_2 [1 - p_1]$$

and (N, I) is an equilibrium if

$$p_2 G > c_2 \text{ and } c_1 > G p_1 [1 - p_2]$$

We can now look at the plane with c_1 and c_2 as its axes, position the other parameters on this and analyze when (I, I) , (N, I) , (I, N) and (N, N) are Nash equilibria. The $c_1 - c_2$ plane is divided into five regions by the above inequalities on c_1 and c_2 . In the lower left region the only possible equilibria are those where both firms choose to invest and in the upper right region the only equilibria are those where neither chooses to invest. Between these regions is one where there are two possible outcomes, (N, I) and (I, N) , and to the upper left the only possible outcomes are (I, N) and to the lower right (N, I) . If both firms are identical then the figure is completely symmetric and of course $c_1 = c_2$ so we are restricted to the diagonal. We therefore have three possible outcomes: (I, I) for low c values, (N, N) for high c values; in between both (N, I) and (I, N) are possible. The asymmetric regions are not possible if the firms are identical.

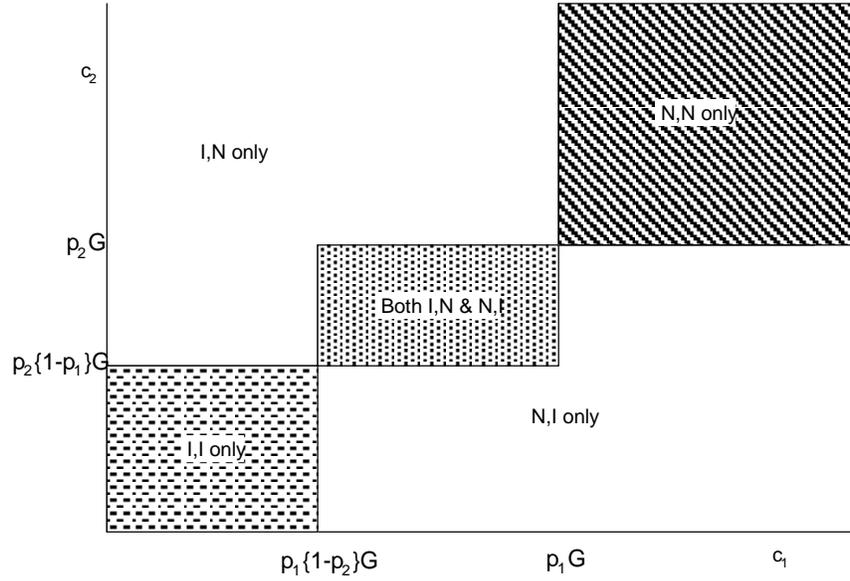


Figure 4: Equilibria of the DR&D game as a function of costs c_2 and c_1 and other parameters.

There are two possible generalizations to an n -agent framework. Scenario 1 is where the firm tries to keep the information proprietary but there is a probability $\tilde{p}_j > 0$ that any other firm obtains the information from firm j if firm j is successful. Scenario 1 is the analog of the computer network case in that once the information becomes public all firms can obtain the information, just as once a virus is spread from one computer it affects all other computers. Scenario 2 is where the information once discovered leaks, but only to one specific firm. Let p_{ji} be the probability of a leak from j to i if j is successful. Scenario 2 is the analog of the airline security case.

We examine Scenario 1 so that each firm undertaking investment provides a positive external effect to the others and in so doing reduces its own incentives to invest. This external effect is referred to as technological spillover in the endogenous growth

literature - see Aghion and Hewitt [1]. In this case the external benefits to firm 1 of four firms consist of terms such as the following where $\{2, 3, 4\}, \{3, 4\}$ etc. now denote the set of firms that are *not* investing in R&D:

$$X_1(4, \{2, 3, 4\}) = 0$$

$$X_1(4, \{3, 4\}) = G\tilde{p}_2$$

$$X_1(4, \{4\}) = G\{\tilde{p}_2 + (1 - \tilde{p}_2)\tilde{p}_3\}$$

$$X_1(4, \emptyset) = G\{\tilde{p}_2 + (1 - \tilde{p}_2)\tilde{p}_3 + (1 - \tilde{p}_2)(1 - \tilde{p}_3)\tilde{p}_4\}$$

If no other firms invest, then the external effects are clearly zero. If firm 2 invests (3 and 4 do not) then the expected gain to one is $G\tilde{p}_2$: if firms 2 and 3 invest then we have the expected gain from firm 2 plus the expected gain from firm 3 conditional on there being no gain from firm 2 etc. This is the same pattern as in the computer security case except that we are now dealing with gains rather than losses.

If firms in the set K are not investing in research then the total expected externality to firm 1 is given by

$$X_1(K) = G \sum_{j \notin K} \tilde{p}_j \prod_{k < j, k \notin K} (1 - \tilde{p}_k) \quad (12)$$

The condition for I to be a dominant strategy when there are n firms with all investing is that

$$c_1 < p_1 [G - X_1(\emptyset)] \equiv c_1(\emptyset) \quad (13)$$

Here as before $c_1(\emptyset)$ is the maximum cost to agent 1 consistent with I being the best choice for 1 when all other firms invest in research. We expect that $X_1(\emptyset)$ increases with n , although unlike the identical agent case we cannot establish a precise limit. If all agents are identical then from K-H [27] $\lim_{n \rightarrow \infty} X_1(\emptyset) = G$.

In the case of three firms we can conduct an analysis similar to the 2-person case of the possible Nash equilibria. For the case of identical companies we find the following pattern of equilibria:

For c between zero and $pG(1 - \tilde{p})^2$ the only equilibrium is where all invest. For $pG(1 - \tilde{p})^2 < c < pG(1 - \tilde{p})$ we have mixed equilibria where two firms invest and one does not - there are three such equilibria and each is equally likely. Again there are mixed equilibria for $pG(1 - \tilde{p}) < cpG$, in this case with one firm investing and two not. Finally for $pG < c$ the only equilibrium is where no one invests.

The 2 and 3-agent pattern of equilibria generalizes straightforwardly to the n -agent case when all agents are identical - this is the case of the diagonal in figure 4.

In the n -agent identical case the regime changes for the Nash equilibria occur at c values given by $pG(1 - \tilde{p})^{n-1}$, $pG(1 - \tilde{p})^{n-2}$, etc. For $c < pG(1 - \tilde{p})^{n-1}$ all invest: for $pG(1 - \tilde{p})^{n-1} < c < pG(1 - \tilde{p})^{n-2}$ all but one invest, for $pG(1 - \tilde{p})^{n-2} < c < pG(1 - \tilde{p})^{n-3}$ all but two invest, and in general for $pG(1 - \tilde{p})^{n-k} < c < pG(1 - \tilde{p})^{n-k-1}$ all but k will invest.

There is an obvious intuition behind this result. The more individuals who invest in R&D, the lower the cost will have to be for you to want to invest yourself.

8.2 Gains for being first

Now we extend the results of the previous section by assuming that there is some advantage to a firm that discovers the information first, even if this eventually percolates to all the others. There could for example be a patent giving property rights for a limited period of time, so that we are studying a patent race - see e.g. [8], [9]. Most literature on patent races assumes that all benefits accrue to the winner - here in contrast we are assuming that some significant benefits accrue to others, perhaps through the ability to licence inventions or to build on them, as in some of the models in Aghion and Hewitt [1]. We formalize this by assuming that the payoff to acquiring the information is F if you are the first to do so, and $G < F$ otherwise. In this case the payoff matrix in the two firm case becomes

	I	N
I	$Y - c_1 + p_1 F + (1 - p_1) \tilde{p}_2 G, Y - c_2 + p_2 F + (1 - p_2) \tilde{p}_1 G$	$Y - c_1 + p_1 F, Y + \tilde{p}_1 G$
N	$Y + \tilde{p}_2 G, Y - c_2 + p_2 F$	Y, Y

and the condition for investing to be a dominant strategy for firm 1 is

$$c_1 < p_1 [F - \tilde{p}_2 G]$$

Not surprisingly the range of costs for which investing can be a dominant strategy is now larger and investment is more likely. In the many firm version of this case the formula (12) still describes the externalities received by firm 1 when firms in K are

not investing and formula (13) becomes

$$c_1 < p_1 [F - X_1(\emptyset)]$$

with G replaced by F . In the identical firm case we have as before that

$$\lim_{n \rightarrow \infty} X_1(\emptyset) = G$$

Since $G < F$ there is still an incentive to invest in R&D even with full spillovers and a large number of firms.

9 Vaccination

The decision facing an individual deciding whether to be vaccinated against an infectious disease is similar to the IDS problem in two respects. Catching these diseases normally conveys immunity so that you can only catch the disease once. In other words damages are non-additive. Secondly, the risk that each person faces depends on whether others are vaccinated - security is interdependent. You can catch the disease from the environment - i.e. from a non-human host - or from another person. If everyone else is vaccinated then the remaining person faces only the risk of catching the disease from a non-human host. Like the R&D problem, there is less of an incentive to adopt the vaccine as more people have protected themselves against the disease.

Assume that it costs c to be vaccinated: this may reflect a combination of cash

costs, psychological costs and possible adverse reactions. If someone catches the disease then the total cost to them is L (for loss). There are non-human hosts for the infectious agent, so that one can be infected even if no one else is. Cholera is a disease of this type: cholera pathogens are resident in the environment even when the disease is not present in humans. The alternative case can be formulated as a special case of this more general situation. Smallpox appears to be a disease that is not endemic in the environment, although a terrorist group could play the role played by non-human hosts in the other case. In the absence of deliberate infection by an enemy, we could not normally catch smallpox unless someone else were already infected. We let p be the probability of catching the disease even if no one else has it: this is the environmental risk of the disease, the background risk (positive for cholera and zero for smallpox). i is the probability that someone who has the disease will infect someone else who is not vaccinated, and ip is the chance of catching the disease and infecting another susceptible person. Y is person i 's initial income or welfare, the reference point from which welfare changes are measured. We denote the product ip by q , as this will occur frequently.

In the two person case we have the following payoff matrix to the strategies of being vaccinated (V) and not being vaccinated (NV):

	V	NV
V	$Y - c, Y - c$	$Y - c, Y - pL$
NV	$Y - pL, Y - c$	$Y - pL - (1 - p)qL, Y - pL - (1 - p)qL$

If both are vaccinated then each has a payoff of $Y - c$, initial income net of the cost

of vaccination. If only one is vaccinated then her payoff is $Y - c$, and the other's is $Y - pL$: the latter person runs no risk of infection from the former as she is vaccinated and by assumption cannot transmit the disease.

In the case in which neither individual chooses to be vaccinated, the payoffs are the initial wealth Y minus the expected losses from two sources: (1) from an infection from the environment pL and (2) from infection by the other person qL , which only matters if you have not already been infected $(1 - p)$. From this payoff matrix it is clear that:

1. When $c < pL$, (V, V) is a Nash equilibrium.
2. For $pL < c < pL + (1 - p)qL$, both (N, V) or (V, N) are equilibria, and
3. For $(1 - p)qL + pL < c$ then (NV, NV) is the equilibrium.

If the cost associated with a vaccination is sufficiently low then both individuals will want to be protected. As c increases then only 1 person will want to be vaccinated. If c is sufficiently high then neither person will want to be protected. This is likely to occur if there is a sufficiently high probability of severe side-effects from the vaccine. The critical values of c at which the equilibrium changes are the expected loss from infection if the other person is vaccinated (pL), and the expected loss from infection if she is not $L(p + (1 - p)q)$. Here $(p + (1 - p)q)$ is the probability of infection if neither is vaccinated. As we shall see below, this structure persists as we consider situations with more people. For a more general discussion of the vaccination problem see Heal and Kunreuther [17].

10 Conclusions

This paper has modeled the management of risks that are discrete and interdependent, and examines how groups of agents react to these risks. The combination of non-additive damages and interdependence of risks gives rise to a novel intellectual structure. This structure is common to a wide range of problems that include airline security, computer network security, bankruptcy and risk-management within an organization, R&D and vaccination. A key aspect of all of these problems is that all agents in the group face the same policy choice, and that the incentives that they have to make this choice depend on the actions of others. The signs of this interdependence vary in the different cases: we have strategic complementarity in the airline and bankruptcy and network cases, and substitutability in the other cases. Lakdawalla and Zanjani [29] also investigate ways in which the public sector can be involved in reducing similar negative externalities. In addition Keohane and Zeckhauser [26] discuss ways of dealing with externalities associated with terrorism when there are threats that affect specific individuals who can contaminate others. They also review collective threats where the number of people exposed to a threat affects the probability of a terrorist attack. This is another version of the problem of endogenous probabilities that we consider in section 5.

Another application of our model, currently topical in the U.S., is to the management of failure risks in electric power grids. Such grids consist of a number of networks transferring power between each other. Each faces a risk that the power transmitted over its network will exceed capacity, causing a local failure that, if not

effectively isolated, can shift power demand to other parts of the network and cause them to be overloaded and fail. In interconnected networks the inability to isolate a failure can cause the initial failure to propagate through the system, which is what happened on 9/3/03. This situation has all the main characteristics of an IDS game: risks are interdependent, the incentive to invest in failure prevention and isolation is reduced if others do not do the same, and damages are non-additive in that once a system has been closed by an overload, additional power surges can do no further damage.

Another issue that has some resemblance to our models is that of bank runs or panics (see Calomiris and Gorton [5], Diamond and Dybvig [10] and Chari and Jaganathan [6]). In these situations the failure of one bank leads depositors to revise their estimates of the safety of remaining banks and can lead to panic withdrawals from these even though they were otherwise facing no risks of failure. This is obviously similar to our concept of "contagion", the spreading of a risk from one agent to another. Via this mechanism one bank behaving in an imprudent manner can impose risks on others even though they themselves behave with the utmost caution.¹⁰ There is however an element of our model that is not present in the models of bank panics just cited: in our model the knowledge that one firm will underinvest in security will reduce the incentives that others have to invest. Hence in equilibrium there may be underinvestment all around. This is not a feature of the banking panic models,

¹⁰For a similar treatment of contagion in the context of insurance see Polonchek and Miller [31] with respect to equity issuances by insurers. In their model an announcement of an issuance by an individual insurers reveals information about the quality of the announcing firm's portfolio as well as the quality of rival firms' portfolios.

although if the banks are aware of the contagion effect then perhaps it should be.¹¹

An interesting feature of the first three cases - airline security, computer network security, bankruptcy and risk-management within an organization - is the possibility of tipping. Tipping occurs when changes in the behavior of a small number of players lead all the rest to change their strategies, thus transforming the equilibrium radically. In such situations, one or a few players are likely to have great leverage over the system as a whole. In our 3-agent numerical example a change of strategy from N to S by one airline leads the other two airlines to also invest in security. Closely associated with the possibility of tipping is that of cascading, where a change of strategy by one agent causes a domino effect that leads a second to change, then a third, and so on until all have changed, a classical "domino effect".

The policy implications are interesting: it may be that the private sector through some coordinating mechanism (e.g. a trade association) or the government can identify those "influentials" or "opinion leaders" whom it is cost-effective to persuade to change their positions. As noted in our example, the tax needed to influence the minimum critical coalition is much less than that needed to influence all players. In [27] we examine private and/or public sector policy interventions that could be used to correct the underinvestment. These include taxes, subsidies, regulations, third party inspections and the use of associations and other coordinating mechanisms.

The equilibria for IDS problems are often inefficient because of the negative ex-

¹¹There is a literature on "fads" and their social transmission which could be thought of as modelling a contagion process -see Bikhchandani Hirshleifer and Welch [3]. However the underlying motivation is rather different from ours - there is no element of risk management or risk spreading in those models.

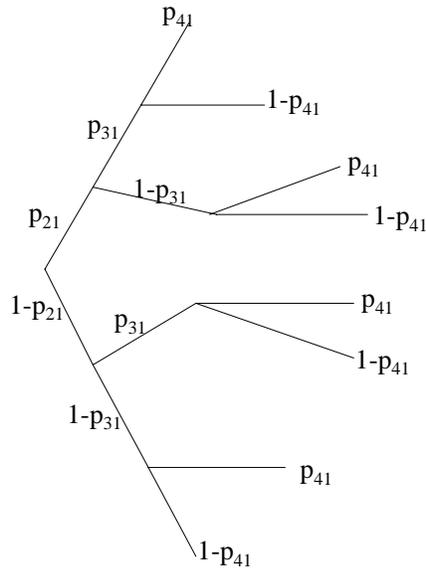


Figure 5: Event tree for the four firm case.

ternal effects between parties. The social return to an investment (in protection, in R&D, in risk management or in infection-prevention) is greater than the private return, thus leading to under-investment. In the special cases in which all agents are identical and the number of agents is very large, we can quantify the under-investment because we have a simple expression for the incentive to invest in security. In the computer network case this incentive approaches zero and in the airline security case it is reduced to about 60% of what its value would be in the absence of external effects.

11 Appendix

The event tree in figure 5 shows the possible ways in which a bag can explode on airline 1 when there are four firms in total and none have invested in security. A

bomb may or may not be transferred from 2, and then may or not be transferred from 3, and then likewise from 4. If 1, 2 or 3 bombs are transferred, the loss is still L . This means that the expected externality imposed on firm 1 is, as stated in section 2

$$X_1(4, \emptyset) = L \{p_{21} + (1 - p_{21})p_{31} + (1 - p_{21})(1 - p_{31})p_{41}\}$$

Proposition 2. Assume that $p_{ij} = \tilde{p}_i \forall j \neq i$. If there is a minimal critical coalition of $k < n$ agents then it must consist of the first k agents ranked by $\Delta_i(\emptyset)$ or equivalently by \tilde{p}_i

Proof. To tip a Nash equilibrium from one at which none invest to one at which all are investing requires that the minimum cost at which investing is justified be raised from below to above the cost of investing for all other than those in the critical coalition. Formally this means that in the initial situation $c_i > c_i(\emptyset) \forall i$ but in the final situation $c_i < c_i(I) \forall i \notin I$ where the agents in the set I are investing in security and form a critical coalition. Now from equation (3) $c_i(\emptyset) = p_{ii}[L - X_i(\emptyset)]$ and $c_i(I) = p_{ii}[L - X_i(I)]$. As agents in I change strategy the changes in the maximum cost consistent with investing in security are $c_i(I) - c_i(\emptyset)$ and to tip the equilibrium it is necessary (and sufficient) that $c_i(I) - c_i(\emptyset) \geq c_i - c_i(\emptyset) \forall i \notin I$. Rank agents by the size of $\Delta_i(\emptyset)$, without loss of generality ordering them so that $\Delta_1(\emptyset) \geq \Delta_2(\emptyset) \geq \Delta_3(\emptyset) \geq \dots$. If agent 1 switches then maximum cost consistent with investing rises by $p_{ii}\Delta_1(\emptyset) \forall i \notin I$. If agents 1 and 2 switch then the returns rise by $p_{ii}(\Delta_1(\emptyset) + \Delta_2(\emptyset))$, etc. Let $\max_{i \notin K} [c_i - c_i(\emptyset)] \leq p_{ii} \left(\sum_{j \leq k} \Delta_j(\emptyset) \right)$ where k is the smallest number for which this holds. Then K is the minimum critical coalition where

$K = \{1, 2, \dots, k\}$ are the first k agents ranked by $\Delta_j(\emptyset)$. ■

Proposition 4 *Under the assumption of the paper, a Nash equilibrium in pure strategies exists for the models of sections 2 and 5.*¹²

Proof. We prove existence of an equilibrium constructively, giving an algorithm which will terminate by locating an equilibrium. We consider the model of section 2: essentially the same argument will apply for the other models considered in the text.

First set all strategies at Y , so that all firms are investing in security. If each firm is playing a best response we have an equilibrium and we are done. Suppose that without loss of generality the first k firms are not picking best responses at this configuration: change their strategies to N . It is clear that for these firms N is a dominant strategy, as when all others are picking S their environment is most conducive to S being the best strategy. If some other firm switches from S to N then this can only make N more attractive to firms from 1 to k : hence N is a dominant strategy for them. Next check whether we have an equilibrium when firms 1 to k choose N and $k + 1$ to n choose S . If yes, we are done.

If not, there are some firms in $k + 1$ to n for which N is the best response to the strategies now being played by the others: change their strategies to N . Now check again if we have a Nash equilibrium. If yes, we are again done. If not, proceed as before: change the strategies of the firms for which S is not a best response to N .

This process will terminate either when all firms are choosing N , which will be a Nash equilibrium, or at a point when there is a Nash equilibrium with some firms

¹² *This argument is taken with grateful thanks from Michael Kearns, personal communication [23].*

choosing N and others choosing S .

Next we extend this argument to the case of endogenous probabilities considered in section 5. For the argument to work in this case we require that it still be the case that a firm is most likely to choose S when all others are also choosing S and that if in such a situation it chooses N then it will always choose N . But this is implied by the assumption of section 5 that the total externality imposed on a firm decreases as the number of other firms investing increases. Hence the same argument applies in the case of endogenous probabilities. ■

References

- [1] Aghion Philippe and Peter W. Hewitt 1997 *Endogenous Growth Theory* MIT Press.
- [2] 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>
- [3] Bikhchandani Sushil, David Hirshleifer and Ivo Welch 1998 "Learning from the Behavior of Others: Conformity, Fads, and Informational Cascades" *The Journal of Economic Perspectives*, Vol. 12, No. 3. (Summer, 1998), pp. 151-170.
- [4] Bulow Jeremy I., John D. Geanakoplos and Paul D. Klemperer 1985. "Multi-market Oligopoly: Strategic Substitutes and Complements". *Journal of Political Economy*, vol 93, no 313, 488-511.

- [5] Calomiris, Charles W. and Gary Gorton 2000. "The Origins of Banking Panics" in Charles W. Calomiris (ed) *U.S. Bank Deregulation in Historical Perspective*, Cambridge University Press.
- [6] Chari, V.V. and Ravi Jaganathan, 1988. "Banking Panics, Information and Rational Expectations Equilibrium." *Journal of Finance* 43: 749-760.
- [7] Crawford, Vincent and Hans Haller, Hans 1990. "Learning How to Cooperate: Optimal Play in Repeated Coordination Games" *Econometrica* Vol 58: Issue 3 (May 1990) 571 - 595.
- [8] Dasgupta Partha and Joseph E. Stiglitz 1981. "Resource Depletion under Technological Uncertainty." *Econometrica* 49, 1: 85-104
- [9] Dasgupta Partha and Joseph E. Stiglitz 1980. "Uncertainty Industrial Structure and the Speed of R&D." *The Bell Journal of Economics* 11, 1: 1-28.
- [10] Diamond, Douglas and Phillip Dybvig 1983. "Bank Runs, Liquidity and Deposit Insurance." *Journal of Political Economy* 91: 401-419.
- [11] Dixit Avinash K. 1988. "A General Model of R & D Competition and Policy." *The RAND Journal of Economics* 19, 3: 317-326.
- [12] Dixit Avinash K 2002. "Clubs with entrapment." Available at www.princeton.edu/~dixitak/home

- [13] Farrell Joseph and Garth Saloner 1985. "Standardization, Compatibility, and Innovation" *The RAND Journal of Economics*, Vol. 16, No. 1. (Spring, 1985), pp. 70-83.
- [14] Gladwell Malcolm 2000. *The Tipping Point* Little Brown and Co.
- [15] Grossman Gene M. and Carl Shapiro 1987. "Dynamic R & D Competition." *The Economic Journal* 97, 386: 372-387.
- [16] Heal Geoffrey 1994. "The Formation of International Environmental Agreements". Pages 301-322 of C. Carraro (ed) *Trade Innovation and the Environment*, Kluwer.
- [17] Heal Geoffrey and Kunreuther Howard 2002. "The Vaccination Game." Working paper, Columbia Business School and Wharton Risk Management and Decision Processes Center, in preparation.
- [18] Heller, Walter 1986 "Coordination Failure in Complete Markets with Applications to Effective Demand". In *Equilibrium Analysis: Essays in Honor of Kenneth J. Arrow Vol II*, ed. W.P.Heller, R.M. Starr and D.A. Starrett, Cambridge University Press., 1986.
- [19] Hirschleifer, J. 1983. "From weakest link to best shot: The voluntary provision of public goods." *Public Choice* 41, 169-186. See also Hirschleifer J. "From weakest link to best shot: correction." *Public Choice* 46, 221-223.
- [20] Hoch Stephen and Howard Kunreuther (ed) 2001 *Wharton on Making Decisions* New York: John Wiley.

- [21] Kaplan Edward H., David L. Craft and Lawrence M. Wein 2002. "Emergency response to a smallpox attack: the case of mass vaccination." *Proceedings of the National Academy of Sciences* August 6 2002 Volume 99 10935-10940.
- [22] Katz Michael and Carl Shapiro 1994. "Systems competition and network effects". *Journal of Economic Perspectives* Vol. 8 No. 2 (Spring 1994): 93-115.
- [23] Kearns, Michael. Personal Communication, March 2003, Department of Computer and Information Science, University of Pennsylvania. www.cis.upenn.edu/~mkearns/
- [24] Kearns Michael and Y. Mansour 2002. "Efficient Nash computation in summarization games with bounded influence." In *Proceedings of the Conference on Uncertainty in Artificial Intelligence*.
- [25] Kearns Michael and Luis Ortiz 2003. "Algorithms for Interdependent Security Games." To appear, *Proceedings of the Seventeenth Conference on Neural Information Processing Systems*.
- [26] Keohane Nathaniel O. and Richard J. Zeckhauser 2002. "The Ecology of Terror Defense." Working paper, Kennedy School of Government, Harvard University. *Journal of Risk and Uncertainty*, forthcoming, Special Issue on Terrorist Risks.
- [27] Kunreuther Howard and Geoffrey Heal 2002. "Interdependent Security: the Case of Identical Agents." Working Paper, Columbia Business School and Wharton Risk Management and Decision Processes Center. *Journal of Risk and Uncertainty*, forthcoming, Special Issue on Terrorist Risks.

- [28] Kunreuther Howard and Geoffrey Heal 2002. "A Firm Can Only Go Bankrupt." Working paper, Columbia Business School and Wharton Risk Management and Decision Processes Center.
- [29] Lakdawalla Darius and George Zanjani 2003 "Insurance, Self Protection and the Economics of Terrorism." Paper Presented at National Bureau of Economic Research Workshop, February 2003.
- [30] Orszag, Peter and Joseph Stiglitz 2002 "Optimal Fire Departments: Evaluating Public Policy in the Face of Externalities" January 2002 Brookings Institution Working Paper.
- [31] Polonchek John and Ronald K. Miller 1999. "Contagion Effects in the Insurance Industry" *Journal of Risk and Insurance* 66, 3: 459-475.
- [32] Sandler, Todd, 2003. Collective Action and Transnational Terrorism. School of International Relations, University of Southern California.
- [33] Sandler, Todd, and Daniel G. Acre 2003. Terrorism and Game Theory. *Simulations and Gaming*, Vol 34(3) September 2003.
- [34] Sandler Todd and Simon Vicary 2001. "Weakest-link public goods: giving in-kind or transferring money in a sequential game." *Economics Letters* 74 (2001) 71-75.
- [35] Schelling, Thomas 1978. *Micromotives and Macrobehavior*. New York: Norton

- [36] Vicary Simon and Todd Sandler 2002. "Weakest-link public goods: Giving in-kind or transferring money." *European Economic Review* 46 (2002) 1501-1520.
- [37] Watts Duncan J. 1999. *Small Worlds*. Princeton University Press.
- [38] Woo, Gordon (2002) "Quantitative Terrorism Risk Assessment", *Journal of Risk Finance*, Vol. 4 (1): 7-14.