

You Only Die Once: Managing Discrete Interdependent Risks

Geoffrey Heal and Howard Kunreuther¹²

December 2002

¹Heal is the Garrett Professor 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 and co-director of the Wharton Risk Management and Decision Processes Center kunreuther@wharton.upenn.edu. He also is a Visiting Research Scholar at the Columbia University Earth Institute.

²`swp35\Docs\securitygeneralAEA.tex`

Abstract

This paper extends our earlier analysis of interdependent security issues to a general class of problems involving discrete interdependent risks. 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 at 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 extend our earlier analysis to cover heterogeneous agents and characterize the tipping phenomenon.

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, and other discrete events. There are other 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 such events, plus 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.

The key point is this: an agent's incentive to invest in risk-reduction measures depends on how he expects the others 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. This situation does not have the structure of a prisoners' dilemma game, even though it has some similarities (see [11]). It contains elements of the coordination problem discussed by Heller [7], Crawford [4] and others.

1.1 Features of the Problem

There are several different versions of this *interdependent security (IDS)* problem and 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. It is not useful in these examples 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 his company will be sent into bankruptcy depends not only on how he manages his 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. The risk depends in part on the behavior of others. The unfavorable outcome is discrete in that it either happens or does not.

1.2 Problem Structure

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. These are the security of airlines and 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 one 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. 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 depend on what others do. Suppose that there are a large number of agents in the system. In K-H [11] 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. 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, 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. Each unit 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 fits in this category. In such a situation the incentive that any unit has 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. In this case the externalities are positive rather than negative .

Our earlier paper [11] 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 show that this introduces the possibility of tipping in the spirit of Schelling [13]. For example 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. 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.

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. We model this phenomenon here.

The next two sections of the paper develop an IDS model with the probabilities and risks differ between agents and then characterizes the structure of the Nash equilibria. Section 4 then considers the case of tipping and shows that those agents which produce the greatest negative externalities occupy a strategic position: 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.

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 identical. Let p_i be the probability that airline i loads a bag on any trip containing a bomb on board at its own facilities in the absence of a security system. Let q_i be the probability that a bag containing a bomb is loaded onto airline i on any trip and is then transferred to another airline, again absent a security system. We expect that $p_i + q_i < 1$ so that there is some chance, $1 - p_i - q_i$, that the airline does not load a bomb. All airlines are equally likely to be recipients of a transferred bomb, so that the chance of any one receiving the bomb from airline i is $q_i / (n - 1)$. Each airline can either invest in a security system S at a cost per trip of c_i or not invest N . Security systems are assumed to be completely effective so that they reduce the chance of a bomb coming through the airline's own facility to zero. In the event that a bomb explodes on a plane the loss is L . The initial income of an airline is Y .

In the case of just two airlines A_1 and A_2 this 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 , q_2L while A_2 has a payoff of income Y minus the expected loss from a bomb loaded on to its plane, p_2L . If neither invests then A_1 has a payoff of income Y minus the expected loss from a bomb loaded on to its own plane p_1L minus the expected loss from a bomb transferred from A_2 , q_2L , conditioned on there being no explosion from a bomb loaded by A_1 itself $(1 - p_1)$. A_2 's payoff is determined in a similar fashion.

A_1/A_2	S	N
S	$Y - c_1, Y - c_2$	$Y - c_1 - q_2L, Y - p_2L$
N	$Y - p_1L, Y - c_2 - q_1L$	$Y - p_1L - (1 - p_1)q_2L, Y - p_2L - (1 - p_2)q_1L$

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

$$c_1 < p_1L \text{ and } c_1 < p_1 [1 - q_2] L \quad (1)$$

The condition that $c_1 < p_1L$ is clearly what we would expect from a single airline operating on its own. The tighter condition that $c_1 < p_1 [1 - q_2] L$ reflects the risk imposed by a firm without security on its partner: 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 three airlines this takes the values

$$\begin{aligned} X_1(3, \{2, 3\}) &= 0 \\ X_1(3, \{3\}) &= Lq_2/2 \\ X_1(3, \emptyset) &= L \{q_2/2 + (1 - q_2/2)q_3\} \end{aligned}$$

The last case reflects the fact that a loss from a bomb transferred from the third airline is possible only if there is no loss from a bomb transferred from the second.

For four firms

$$\begin{aligned} X_1(4, \{2, 3, 4\}) &= 0 \\ X_1(4, \{3, 4\}) &= Lq_2/3 \\ X_1(4, \{4\}) &= L \{q_2/3 + (1 - q_2/3)q_3\} \\ X_1(4, \emptyset) &= L \{q_2/3 + (1 - q_2/3)q_3/3 + (1 - q_2/3)(1 - q_3/3)q_4/3\} \end{aligned}$$

In all cases when transfers from more than one airline are possible then the losses from the second and subsequent firms have to be conditional on there being no losses from previous transfers.¹

¹We are using a particular ordering of the other firms, the natural order, but it can be shown that the order in which we list firms in $\{K\}$ does not affect the final value of $X_i(n, \{K\})$.

For n airlines when none of them invest in security this generalizes to the following formula for the externality inflicted on the first:

$$X_1(n, \emptyset) = L \sum_{j=2}^{j=n} \frac{q_j}{n-1} \prod_{k=2}^{k=j-1} \left(1 - \frac{q_k}{n-1}\right)$$

where it is understood that $\prod_{k=2}^{k=j-1} \left(1 - \frac{q_k}{n-1}\right) = 1$ when $j = 2$. Again this result is independent of the order in which we consider the other firms. If firms in the set $\{K\}$ are investing in security then the total externality to firm 1 is given by

$$X_1(n, \{K\}) = L \sum_{j \notin \{K\}} \frac{q_j}{n-1} \prod_{k < j, k \notin \{K\}} \left(1 - \frac{q_k}{n-1}\right) \quad (2)$$

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

$$c_1 < p_1 [L - X_1(n, \emptyset)] = c_1(n, \emptyset) \quad (3)$$

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 [11] that if all agents are identical then $\lim_{n \rightarrow \infty} X_1(n, \emptyset) = L(1 - e^{-q})$.

3 The Structure of Nash Equilibria

The nature of the Nash equilibrium in the interdependent security model naturally depends on the parameters. For convenience we repeat the two-person payoff matrix:

A_1/A_2	S	N
S	$Y - c_1, Y - c_2$	$Y - c_1 - q_2L, Y - p_2L$
N	$Y - p_1L, Y - c_2 - q_1L$	$Y - p_1L - (1 - p_1)q_2L, Y - p_2L - (1 - p_2)q_1L$

From this it is clear that (S, S) is a Nash equilibrium if $c_i < p_iL$. and is a dominant strategy if $c_i < p_iL(1 - q_j)$ where i and j are 1 or 2. (N, N) is a Nash equilibrium if $c_i > p_iL(1 - q_j)$ and a dominant strategy if $c_i > p_iL$. From these inequalities we note

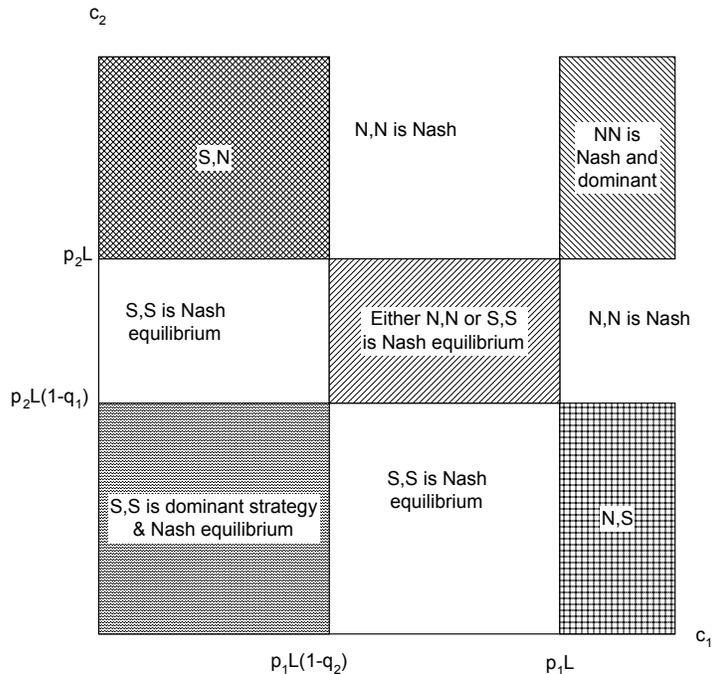


Figure 1: Nash equilibria as a function of c_1 and c_2 .

that (S, S) and (N, N) are both Nash equilibria if $p_i L (1 - q_j) < c_i < p_i L$. Finally if $c_1 > p_1 L$ but $c_2 < p_2 L (1 - q_1)$ (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 [11].

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 $c_i(n, \emptyset)$ is the maximum cost at which agent i will invest in security if no others are investing. Clearly if $c_i > c_i(n, \emptyset) \forall i$ 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 $E \subset \{n\}$ of agents choosing S such that $c_i < c_i(n, \{E\}) \forall i \in E$ and $c_i > c_i(n, \{E\}) \forall i \notin E$.

4 Nash Equilibria & Tipping

For cases where there is more than one Nash equilibrium, 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 spirit of both Schelling ([13]) and Gladwell ([5]). 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 $\{N, N, \dots, N\}$ and specify the conditions for a subset of these agents to lead all the others to invest in security producing an equilibrium $\{S, S, \dots, S\}$

Consider a Nash equilibrium where all agents choose strategy N . The minimal critical coalition $\{K\}$ is the smallest group of firms that by switching from N to S can tip the equilibrium to one where all firms invest. Formally consider a Nash equilibrium such that $c_i > c_i(n, \emptyset) \forall i$, so no agents invest in security, and $c_i < c_i(n, \{K\}) \forall i \notin K$. The firms in K form a critical coalition in that if they switch from N to S all other firms will follow suit. If a set $\{K\}$ has this property and no proper subset has it, then that set is a minimum critical coalition, MCC.

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 $\{K\}$ are investing in security as the change in the total externality to j when i switches from not investing to investing. Denote this by $E_{ij}(\{K\})$. Formally this is

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

From the definitions of $X_j(n, \{K\}) - X_j(n, \{K + i\})$ we have:

$$E_{ij}(\{K\}) = L \frac{q_i}{n-1} \prod_{k \notin \{K\}} \left(1 - \frac{q_k}{n-1}\right) \quad (4)$$

From equation (4) we see that this externality generated by i does not depend on the index j , as it does not appear on the right hand side of (4). The reason for this is that we have assumed in this model that the probability of a bag being transferred from i to j is the same for all j , as is the loss incurred in the event that a bag explodes. From this it follows that the total externality generated by i is just $(n-1)$ times (4):

$$E_i(\{K\}) = (n-1) E_{ij}(\{K\}) = L q_i \prod_{k \notin \{K\}} \left(1 - \frac{q_k}{n-1}\right) \quad (5)$$

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 probability of a transfer from i conditioned on the firms not in 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 $E_i(\emptyset)$.

The next issue is: are the firms with the highest values of $E_i \{\emptyset\}$ the ones which will “tip” a Nash equilibrium from not investing to investing? Is $E_i \{\emptyset\}$ a good measure of a firm’s “leverage”? The following proposition shows that it is.

Proposition 3 *If there is a minimal critical coalition of $k < n$ agents then it must consist of the first k agents ranked by $E_i \{\emptyset\}$.²*

Proof. The proof is in the Appendix. ■

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. $p_i L$) minus the likelihood that the agent will be harmed by unprotected agents multiplied by the resulting loss L . Order the unprotected agents so that agent 1 has $\max\{E_i \{\emptyset\}\}$, $i \in n \setminus K$, 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 MMC.

The following numerical example demonstrates that such an MMC 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 $q_1 = q_2 = 0.1$ and $q_3 = 0.5$ and $L = 1000$. In addition $p_1 = p_2 = 0.1$ and $c_1 = c_2 = 85$. We do not specify p_3 or c_3 . In this setting

$$c_1(\emptyset) = p_1 \left[1000 - \frac{0.1 \times 1000}{2} - \left(1 - \frac{0.1}{2} \right) \frac{1000}{2} \right] = 71.25$$

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 for any cost above 75 not investing is a dominant strategy for both the other firms.

Now suppose that for some reason airline 3 changes policy and invests. It now imposes no externality on the other firms and so can be left out of the analysis of their decisions. To understand the choices of firms 1 and 2 we simply have to apply inequality (1) in the present case, 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

²Ties will be broken randomly if needed.

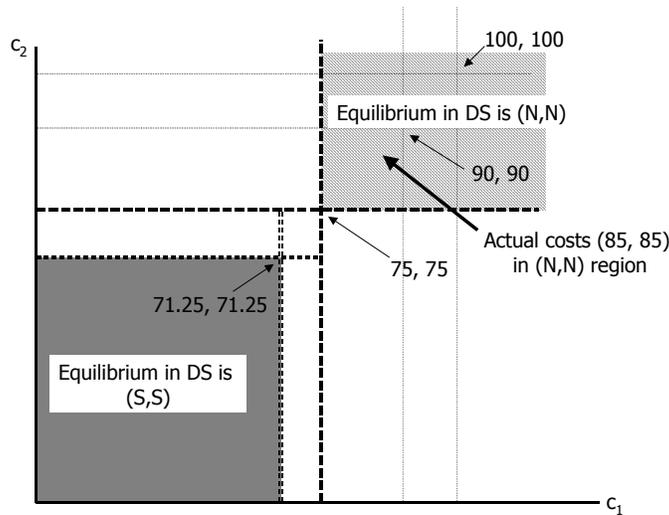


Figure 2: Equilibria for firms 1 and 2 when 3 does not invest and imposes externalities on them. In this case $(85, 85)$ is in the region in which not investing is a dominant strategy.

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 their 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. And it is easy to verify that airline 3 imposes the largest externalities on the other airlines in accordance with Proposition 1 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 to understand 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.³ 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. 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

³We 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.

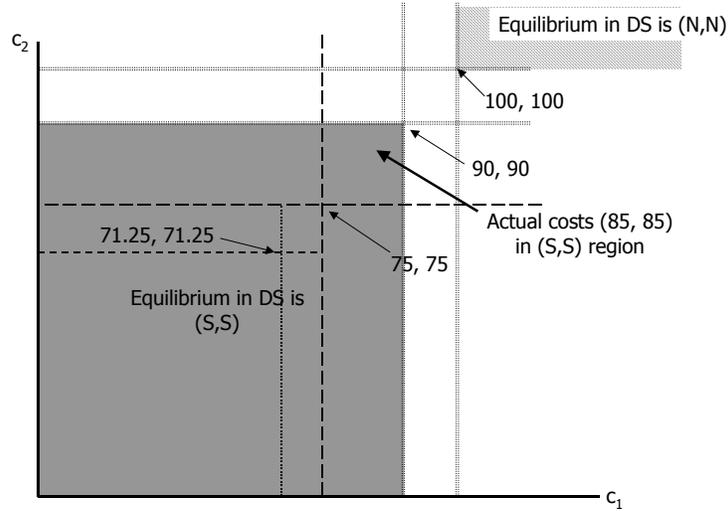


Figure 3: Equilibria for firms 1 and 2 when 3 invests and imposes no externalities on them. In this case (85, 85) is in the region in which investing is a dominant strategy.

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.

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. This is similar to the problem of theft protection. If a house announces that it has installed an alarm, then burglars are likely to turn to others houses as targets. ([11]) In the case of airline security, terrorists are more likely to focus on targets which are not as well protected as others, so that the p 's and q 's depend on the investment in security. In our example we assume that the risk faced by an airline that does not invest 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 the relevant probabilities of those not investing in security, $p_i(\#\{K\})$ and $q_i(\#\{K\})$, are increasing in $\#\{K\}$ where $\{K\}$ is the set of airlines investing in security and $\#\{K\}$ is the number of airlines in $\{K\}$. For airlines that have invested in security the ps and qs 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_i and q_i are exogenous:

$$c_i < p_i [L - X_i(n, \emptyset)] = c_i(n, \emptyset)$$

As an increasing number of other firms invest in security, then the probability

p_i will increase as will the value of q_j for all unprotected agents. Hence the value of $X_i(n, \{K\})$ will increase. The expression for the critical value of the cost of investment for airline i is thus :

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

The right hand side of (6) increases in $\#\{K\}$ via p_i but also depends on $\#\{K\}$ through the q_s that enter into the expression for the externality imposed on i . 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 $\sum_{j \notin \{K\}} \frac{q_j(\#\{K\})}{n-1}$ but will also decrease

$\prod_{k < j, k \notin \{K\}} \left(1 - \frac{q_k(\#\{K\})}{n-1} \right)$. 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 the number of firms investing $\#\{K\}$ will increase the right hand side of (6). and raise the value of $c_i(n, \{K\})$. 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 results also implies the above analysis with constant probabilities is qualitatively valid. Of course the computation of $c_i(n, \{K\})$ 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 enemy 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, Proposition 1 on tipping is 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 not invested in security will become a more likely target. It is therefore likely that the endogeneity of probabilities will reduce the size of an MCC.

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 but where the definitions of p_i and q_i may differ. To keep the analysis simple the probabilities are assumed to be exogenous. The qualitative results for each of these problems is similar to the airline security case when the probabilities are endogenous.

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 damaging only one of them (Anderson [2]). Let p_i be the probability that computer i is infected by a virus and q_i be the probability that it

is infected by a virus and then is transmitted to all other computers. Clearly $q_i \leq p_i$ and p and q do not refer to independent events. The other notation is the same as in the airline security problem. Since a computer virus can damage all computers on the network, the probability that computer i infects computer j is now q_i rather than $q_i/(n-1)$ as in the airline problem. So the stochastic negative externalities for the case of four computers are given by:

$$\begin{aligned} X_1(4, \{2, 3, 4\}) &= 0 \\ X_1(4, \{3, 4\}) &= Lq_2 \\ X_1(4, \{4\}) &= L\{q_2 + (1 - q_2)q_3\} \\ X_1(4, \emptyset) &= L\{q_2 + (1 - q_2)q_3 + (1 - q_2)(1 - q_3)q_4\} \end{aligned}$$

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:

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

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

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

$$X_1(n, \{K\}) = L \sum_{j \notin \{K\}} q_j \prod_{k < j, k \notin \{K\}} (1 - q_k) \quad (8)$$

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

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

Here as before $c_1(n, \emptyset)$ is the maximum cost to agent 1 consistent with S being a Nash equilibrium when no other agents invest in security. In K-H ([11]) we show that if all agents are identical the term $X_1(n, \emptyset)$ 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 $X_1(n, \emptyset)$ goes to L as $n \rightarrow \infty$ so that it is again the case that with more agents it is less likely that an agent will choose to invest.

The definitions of Nash equilibrium and minimum critical coalition carry over unchanged from the previous case. It is again the case that agents can be ranked by the value of the externalities that they impose on others when they switch policy so that we can prove an exact analog of proposition 1 for the computer network case:

Proposition 4 2 *A minimal critical coalition of $k < n$ agents must consist of the first k agents ranked by $E_j \{\emptyset\}$.*

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. In the introduction we cited several examples of this type. 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 destroy the company and indeed they did.⁴ 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 these cases the situation is analytically similar to the computer security problem. An organization consists of a group of divisions i , 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 q_i is 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 $q_i \leq p_i$ and p_i and q_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 .⁵ 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 Barings, and Anderson's Houston branch similarly imposed risks on all of Anderson. And as before this 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

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

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

⁴For a more detailed description of the factors causing the collapse of Barings Bank see Hoch and Kunreuther (Chap.1) (2001)

⁵One could also think of c_i as an opportunity cost resulting from avoiding certain deals that might compromise the entire enterprise.

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. For a more extensive analysis of the managerial and organizational implications of this problem, see [12].

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). 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 the solution may be available to some or all others. In this example information has some of the characteristics of a public good and so may become available to others besides the immediate discoverer 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. There is however one important difference. In this case investment decisions are not mutually reinforcing. In the airline security, 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. 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 these organizations also increases. Therefore it is more efficient for them to wait until one of these entities 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 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 q_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

	I	N
I	$Y - c_1 + p_1 G + (1 - p_1) q_2 G, Y - c_2 + p_2 G + (1 - p_2) q_1 G$	$Y - c_1 + p_1 G, Y + q_1 G$
N	$Y + q_2 G, Y - c_2 + p_2 G$	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 one firm invests and the

other 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 $Y + q_1 G$, income plus the expected gain as the information is transferred to it from the successful investor. Finally if both invest then each has a payoff of $Y - c_1 + p_1 G + (1 - p_1) q_2 G$, 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

$$c_1 < p_1 [1 - q_2] G \tag{10}$$

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 $c_1 < p_1 G$. The term $[1 - q_2]$ represents what was previously labeled contagion in ([11]). 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 2 is investing will reduce the incentive that firm 1 has to do likewise.

8.1 Nash Equilibrium for 2-Agent Problem

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

$$c_1 < p_1 [1 - q_2] G \text{ and } c_2 < p_2 [1 - q_1] G$$

Similarly (I, N) is a Nash equilibrium if

$$p_1 G > c_1 \text{ and } c_2 > G p_2 [1 - q_1]$$

and (N, I) is an equilibrium if

$$p_2 G > c_2 \text{ and } c_1 > G p_1 [1 - q_2]$$

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

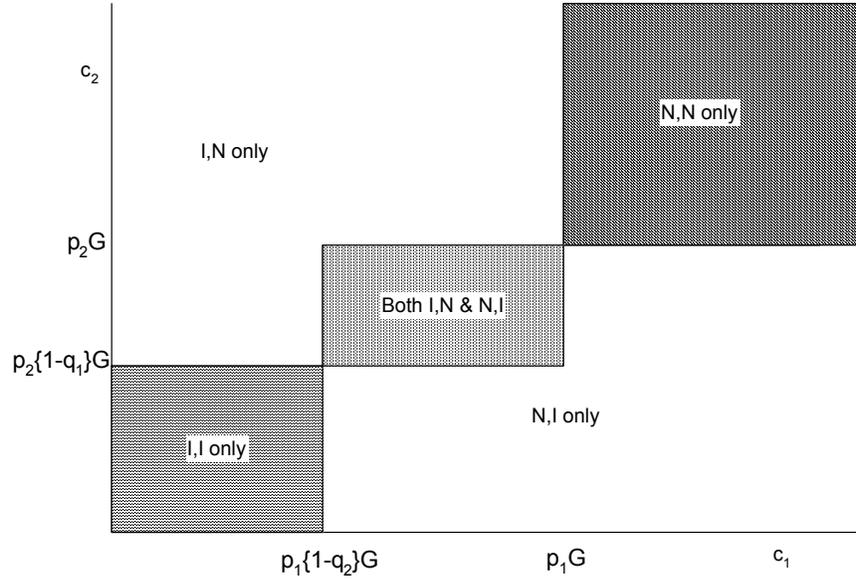


Figure 4: Equilibria of the R&D game as a function of costs c_1 and c_2 and other parameters.

(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. In the case of three firms we can conduct a similar analysis of the possible Nash equilibria. For the case of identical companies we find the following pattern of equilibria:

For c between zero and $pG(1-q)^2$ the only equilibrium is where all invest. For $pG(1-q)^2 < c < pG(1-q)$ 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 $pG(1-q) < cpG$, in this case with one firm investing and two not. Finally for $pG < c$ the only equilibrium is where no one invests.

There are two possible generalizations of this 2 agent problem to an n -agent framework. Scenario 1 is where the firm tries to keep the information proprietary but there is a probability $q_j > 0$ 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. Suppose it is equally likely to leak to any firm: then if q_j is the probability that j makes a discovery and that it leaks, we have $q_j/(n-1)$ as the probability that each firm has access to the discovery of a successful firm via a leak. Scenario 2 is the analog of the airline security case in that there is a $\frac{1}{n-1}$ chance of any firm obtaining the discovery

from another firm just as there is a $\frac{1}{n-1}$ chance of any airline obtaining a bomb transferred from another airline.

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 is as follows where $\{2, 3, 4\}$, $\{3, 4\}$ etc. denotes the set of firms that are *not* investing in R&D):

$$\begin{aligned} X_1(4, \{2, 3, 4\}) &= 0 \\ X_1(4, \{3, 4\}) &= Gq_2 \\ X_1(4, \{4\}) &= G\{q_2 + (1 - q_2)q_3\} \\ X_1(4, \emptyset) &= G\{q_2 + (1 - q_2)q_3 + (1 - q_2)(1 - q_3)q_4\} \end{aligned}$$

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 Gq_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

$$X_1(n, \{K\}) = G \sum_{j \notin \{K\}} q_j \prod_{k < j, k \notin \{K\}} (1 - q_k) \quad (12)$$

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

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

Here as before $c_1(n, \emptyset)$ 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 $X_1(n, \emptyset)$ increases with n , although unlike the identical agent case we cannot establish a precise limit. If all agents are identical then from K-H ([11]) $\lim_{n \rightarrow \infty} X_1(n, \emptyset) = G$.

With respect to a Nash equilibrium the 2-agent pattern of equilibria generalizes to the n -agent case with the regime changes being at c values given by $pG(1 - q)^{n-1}$, $pG(1 - q)^{n-2}$, etc. For $c < pG(1 - q)^{n-1}$ all invest: for $pG(1 - q)^{n-1} < c < pG(1 - q)^{n-2}$ all but one invest, etc.

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. 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>I</i>	<i>N</i>
<i>I</i>	$Y - c_1 + p_1 F + (1 - p_1) q_2 G, Y - c_2 + p_2 F + (1 - p_2) q_1 G$	$Y - c_1 + p_1 F, Y + q_1 G$
<i>N</i>	$Y + q_2 G, Y - c_2 + p_2 F$	Y, Y

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

$$c_1 < p_1 [F - q_2 G] \text{ or } c_1 < p_1 \left[\frac{F}{G} - q_2 \right] G$$

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

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

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

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

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 or you die. 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 a lower 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 in the second category, a disease that is not endemic in the environment. 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.

9.1 The two person case

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 $(1 - p)qL < c$ 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.

9.2 The three person case

For three people the possible strategy combinations are (V, V, V) , (V, V, N) , (V, N, N) and (N, N, N) and all possible permutations of the second and third cases. Payoffs depend only on the total number of strategies of each type and not on their distribution over people and are given by the following matrix, where the first column gives

the set of strategies and the second and third give the payoffs to being vaccinated and to not being vaccinated respectively.

Payoff	V	NV
(V, V, V)	$Y - c$	
(V, V, NV)	$Y - c$	$Y - pL$
(V, NV, NV)	$Y - c$	$Y - pL - (1 - p)qL$
(NV, NV, NV)		$Y - L \left\{ \begin{array}{l} p + (1 - p) \{q + (1 - q)q\} + \\ (1 - p)(1 - q - (1 - q)q) \{qi + (1 - qi)qi\} \end{array} \right\}$

In every case the payoff to being vaccinated V is the same, $Y - c$.

If two people are vaccinated then the other person has a risk p of being infected and a payoff associated with not being vaccinated of $Y - pL$: if one is vaccinated then the risk of infection for a non-vaccinated person is $p + (1 - p)q$, the risk of infection from the environment plus the risk of being infected by the other non-vaccinated person given that he is not already infected by the environment.

If no one is vaccinated then the risk to one of the three of infection is

$$p + (1 - p) \{q + (1 - q)q\} + (1 - p)(1 - q - (1 - q)q) \{qi + (1 - qi)qi\} \equiv TR(3, 0)$$

which is reflected in the final row of the payoff matrix and denoted for brevity $TR(3, 0)$ which is the total risk to one of three people of being infected when none is vaccinated. This is derived as follows:

Individual 1 may be infected by the non-human source (p); with probability $(1 - p)$ this may not happen in which case he may be infected directly via 2 with probability q ; if individual 1 is not infected by 2 he may be infected by 3 which occurs with probability $(1 - q)q$. So $p + (1 - p) \{q + (1 - q)q\}$ give the risk to one of three susceptible people of infection from the background or directly from one of the other two.

It is also possible that one of the three individuals is infected indirectly by one of the other two via three person chains - for example 2 infects 3 who infects 1. The chance of this happening has to be conditioned on not being infected directly by the environment $(1 - p)$ and on not being infected directly by 2 or 3 either $(1 - q - (1 - q)q)$. There are two such three person chains: 2-3-1 or 3-2-1 each with probability qi and the latter has to be conditioned on the former not occurring. These three person chains therefore contribute the term

$$(1 - p)(1 - q - (1 - q)q) \{qi + (1 - qi)qi\}.$$

The Nash equilibria in this case are as follows:

1. $c < pL$: the only equilibrium is (V, V, V) i.e. all are vaccinated.
2. $pL < c < pL + (1 - p)qL$: all equilibria have two people vaccinated and one not, and so are of the form (V, V, NV)
3. $pL + (1 - p)qL < c < TR(3, 0)L$: all equilibria have one person vaccinated and two individuals who are not, and so are of the form (V, NV, NV) .

4. $TR(3, 0) L < c$: the only equilibria have all people unvaccinated, (NV, NV, NV) .

The vaccination problem is studied in more detail in Heal and Kunreuther ([6]). The point of this overview is to show that it is a type of IDS problem having both non-additive damages and interdependent risks where one is less likely to want to protect oneself when others have taken these steps. The solution thus has much in common with the R&D problem.

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.

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 in the sense of Schelling [13]. 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. The policy implications are interesting: the private sector through some coordinating mechanism (e.g. a trade association) or the government can identify those “influentials” or “opinion leaders” where 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.

The equilibria for IDS problems are often inefficient because of the negative external effects between parties. In all cases 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.

K-H (2002) examines various private and/or public sector policy interventions that might be used to correct the under-investment. These included taxes, subsidies, regulations, third party inspections and the use of associations and other coordinating mechanisms. Similar issues are discussed by Keohane and Zeckhauser [10].

11 Appendix

Proof. To tip a Nash equilibrium from one at which none invest to one at which all are investing requires that the returns from investing 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 $c_i(\emptyset) = p_i [L - X_i(\emptyset)]$, $c_i(\{I\}) = p_i [L - X_i(\{I\})]$. As agents in $\{I\}$ change strategy the changes in the returns to 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 $E_i(\emptyset)$, without loss of generality ordering them so that $E_1(\emptyset) \geq E_2(\emptyset) \geq E_3(\emptyset) \geq \dots$. If agent 1 switches then the return to investing rises by $p_1 E_1(\emptyset) \forall i \notin \{I\}$. If agents 1 and 2 switch then the returns rise by $p_i (E_1(\emptyset) + E_2(\emptyset))$, etc. Let $\max_{i>k} [c_i - c_i(\emptyset)] \leq p_i \left(\sum_{j \leq k} E_j(\emptyset) \right)$ 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 $E_j(\emptyset)$. ■

References

- [1] Aghion Philippe and Peter W. Hewitt *Endogenous Growth Theory* MIT Press, 1997
- [2] Anderson, Ross. "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] Barrett, Scott. *Global Disease Eradication*. Working Paper, School of Advanced International Studies, Johns Hopkins University.
- [4] 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.
- [5] Gladwell Malcolm *The Tipping Point* 2000 Little Brown and Co.
- [6] Heal Geoffrey and Kunreuther Howard. "The Vaccination Game." Working paper, Columbia Business School and Wharton Risk Management and Decision Processes Center.
- [7] 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.
- [8] Kaplan Edward H., David L. Craft and Lawrence M. Wein. "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.

- [9] Kaplan Edward H., David L. Craft and Lawrence M. Wein. “Analyzing Bioterror Response Logistics: The Case of Smallpox.”
- [10] Keohane Nathaniel O. and Richard J. Zeckhauser. “The Ecology of Terror Defense.” Working paper, Kennedy School of Government, Harvard University. *Journal of Risk and Uncertainty*, forthcoming, Special Issue on Terrorist Risks.
- [11] Kunreuther Howard and Geoffrey Heal. “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.
- [12] Kunreuther Howard and Geoffrey Heal. “A Firm Can Only Go Bankrupt.” Working paper, Columbia Business School and Wharton Risk Management and Decision Processes Center.
- [13] Schelling, Thomas. (1978) *Micromotives and Macrobehavior*. New York: Norton