You Only Die Once: Managing Discrete Interdependent Risks

Geoffrey Heal and Howard Kunreuther*

April 2002
Very Preliminary Draft

Abstract

This paper extends our earlier analysis of interdependent security issues to a general class of problems involving discrete interdependent risks. These are problems where there is a threat of an event that can only happen once, and the risk of this threat depends on actions taken by others. In such cases any agent's incentive to invest in reducing the threat depends on the actions of others. Security problems at airlines and in computer networks come into this category, as do as do problems of risk management at organizations facing the possibility of bankruptcy. 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

There are certain bad events that can only occur once. Death is the obvious example: an individual's death is irreversible and unrepeatable. More mundane examples are bankruptcy, being struck off a professional register, and other discrete events. In addition there are other events that can in principle occur twice but that are so unlikely and/or so dreadful that one occurrence is all that can reasonably be considered. The events of 9/11/01 are perhaps of this type. A nuclear meltdown in a highly populated region is another. The fact that such events are typically probabilistic, taken together with the fact that the risk that one agent faces is often determined in part by the behavior of others, gives a unique and hitherto unnoticed structure to the incentives that agents face to reduce their exposures to these risks.

*Heal is at the Graduate School of Business and Columbia University - gmh1@columbia.edu and www.gsb.columbia.edu/faculty/gheal. Kunreuther is Visiting Scholar at the Columbia Earth Institute and at the Wharton School of Business of the University of Pennsylvania, kunreuther@wharton.upenn.edu

The key point is that the incentive that any agent has to invest in risk-reduction measures depends on how he expects the others to behave in this respect. If he thinks that they will not invest in security, then this reduces the incentive for him to do so. On the other hand should he believe that they will invest in security, then it may be best for him to do so too. 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. Yet this situation does not have the structure of a prisoners' dilemma game, even though it has some similarities.

We explored the simplest structure of this problem in our earlier paper "Interdependent Security: The Case of Identical Agents" [Kunreuther-Heal (K-H) [2]]. The fundamental question that motivates our research is: "Do organizations facing discrete and interdependent risks, such as airline companies and computer network managers, invest in risk management to a degree that is adequate from either a private or social perspective?" In general the answer is no. The natural next question is of course: "So what should we do about this?"

1.1 Features of the Problem

There are several different versions of this 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 lawyer is disqualified or not. A plane crashes or 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. In contrast to the standard prisoners' dilemma paradigm where the outcomes are specified with certainty, the interdependent security problem involves chance events. 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 as we shall see sufficient to ensure that there can be equilibria at which there is underinvestment in risk-prevention measures. The precise degree of underinvestment depends on the nature of the problem. We focus mainly on the two extremes that span the spectrum of possibilities. These are the security of airlines and 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 [2] 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 problems with discrete and interdependent risks. There are other areas where the same structure arises. Recently Arthur Anderson was sent into bankruptcy by the actions of its Houston branch. Several years ago Barings was likewise destroyed by the actions of a single trader in its Singapore branch. In each case we had multi-unit organizations in which the risk of bankruptcy (a discrete event) faced by any unit was affected by its own choices and by the choices made by other units. In such a situation the incentive that any unit has to reduce bankruptcy risks is reduced by the knowledge that others are not doing so. 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. In section 6 we show that some decisions about 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.

In our earlier paper we studied this problem in the case where all agents are identical and established that the interdependence of risks, together with the discreteness of the payoff, means that the incentive that an agent has to invest in security precautions is reduced, in some cases to zero. Here we extend the analysis to the more general case of agents who differ in important characteristics and show that this introduces an interesting new possibility. This is the possibility of tipping as described by Schelling [3]. In some cases 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 the agents with 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.

2 The Model

Initially we think in terms of the security of airlines, as this is an example that is both topical and canonical. We assume that there are $n \geq 2$ separate airlines. Let p_i be the probability that airline i loads a bag containing a bomb on board at its own facilities in the absence of a security system. q_i is the probability that a bag containing a bomb is loaded onto airline i then transferred to another airline. Clearly we expect that $p_i + q_i < 1$ so that there is some change that the airline does not load a bomb - this chance is $1 - p_i - q_i$. We assume that all airlines are equally likely to be recipients of a transferred bomb, so that the chance of any one receiving the bomb is $q_i/(n-1)$. Each airline can either invest in a security system S 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. Clearly the probability of loading a bag containing a bomb, in the absence of security measures, is $p + q$. The cost of a security system is c_i. 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 invest in security systems then their payoffs 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. If neither invests then A_1 has a payoff of income Y minus the expected loss from a bomb on loaded itself p_1L minus the expected loss from a bomb transferred A_2, q_2L, conditioned on there being no explosion from a bomb loaded by A_1 itself $(1 - p_1)$.

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

$$c_1 < p_1L \text{ and } c_1 < p_1\,[1-q_2]\,L \tag{1}$$

The condition that $c_1 < p_1L$ is clearly what we would expect from a single airline on its own, and 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 insecure airline to the other. This externality plays a critical role in our analysis and we need to understand its structure.

Let $X_i\,(n, \{J\})$ be the expected externality from all other agents to agent i when agents in the set of agents $\{J\}$ invest in security and there are n agents in total. For three firms 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 agents 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 firm are possible then the losses from the second and subsequent firms have to be conditional on there being no losses from previous transfers. We 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 $\{J\}$ does not affect the final value of $X_i\,(n, \{J\})$.

For n firms when none 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

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

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

$$c_1 < p_1\left[L - X_1(n,\emptyset)\right] \equiv c_1(n,\emptyset) \tag{3}$$

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 consistent with agent i investing in security when there are n agents and those in the set $\{K\}$ have already invested in security.

Clearly $X_1(n,\emptyset) > X_1(n,\{2\}) > X_1(n,\{2,3\}) > ... > X_1(n,\{2,3,4,..n-1\})$ so that $c_1(n,\emptyset) < c_1(n,\{2\}) < .... < c_1(n,\{2,3,4...n-1\})$. 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 [2] that if all agents are identical then $\lim_{n\to\infty} X_1(n,\emptyset) = L(1-e^{-q})$.

3 Nash Equilibria & Tipping

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 $\{N,N,....N\}$ is a Nash equilibrium. More generally we can characterize a Nash equilibrium as follows:

Definition 2 A Nash equilibrium is a set $Y \subset \{n\}$ of agents choosing S such that $c_i < c_i(n,\{Y\})\,\forall i \in Y$ and $c_i > c_i(n,\{Y\})\,\forall i \notin Y$.

Clearly from this definition a mixed equilibrium is possible, that is an equilibrium at which some agents invest while others do not. This is a contrast with the identical agent case discussed in our earlier paper K-H 2001.

Definition 3 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.

Firms in a minimum critical coalition are an important group: if they change from not investing to investing then all others follow suit, because by changing they reduce the externalities on others and so reduce the maximum cost of investing at which agents will invest. This can tip other agents from not investing to investing. We will see that there is a rather natural and intuitive characterization of such firms: they are the groups that have the greatest impacts on the decisions made by others. To understand the impact that an agent has on others we need to know the externalities that each agent generates in total on all others, and we turn to this next.

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

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

This would clearly be the same (except for sign) if we reversed the process and had j go from investing to not investing. The total externality from j denoted $E_j\{K\}$ when firms in $\{K\}$ are investing is

$$E_j\{K\} = \sum_i E_{ij}(\{K\}) \tag{4}$$

>From the definitions of E_j and X_i we have:

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

This is the total externality from i to j but it is clear from the formula that it does not in fact depend on the index j. This does not appear on the right hand side of (5). It is clear intuitively that the externality from i to j depends on the probability of a bag being transferred from i to j, which we have assumed to be the same for all j, and also on the loss incurred by j in the event that a bag explodes, which we have also assumed to be the same for all j. From this it follows that the total externality generated by i is just $(n-1)$ times (5):

$$E_i\{K\} = (n-1)E_{ij}(\{K\}) = Lq_i \prod_{k \notin \{K\}} \left(1 - \frac{q_k}{n-1}\right) \tag{6}$$

This is a reasonably compact and intuitive expression for the total externality from i when $\{K\}$ are investing in security. It is the loss times the probability of a transfer from i conditioned on the firms not in K not having inflicted damage on a firm already. We show below that it is relevant in several cases to rank firms by this number $E_i\{K\}$. Specifically we will probably be most interested in the case when no firms are investing in security, so in this case the index is $E_j\{\emptyset\}$.

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 $E_i\{K\}$ a good measure of a firm's "leverage"? The following propositions show that it is.

Proposition 4 A minimal critical coalition of $k < n$ agents must consist of the first k agents ranked by $E_j\{\emptyset\}$.[1]

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. 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. 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 $c_i(\{I\}) - c_i(\emptyset) \geq c_i - c_i(\emptyset)$. Rank agents by the size of $E_i(\emptyset)$, without loss of generality ordering them so that $E_1(\emptyset) \geq E_2(\emptyset) \geq E_3(\emptyset) \geq ......$ If agent 1 switches then the return to investing rises by $p_i E_i(\emptyset)\,\forall i \notin \{I\}$. If agents 1 and 2 switch then the returns rise by $p_i(E_1(\emptyset) + E_2(\emptyset))$, etc. Let $\max_i[c_i - c_i(\emptyset)] \leq p_i \sum_{j \leq k} E_j(\emptyset)$ where k is the smallest number for which this holds. Then $\{K\}$ is the minimum critical coalition where $\{K\} = \{1, 2, ...k\}$ are the first k agents ranked by $E_j(\emptyset)$. ■

Proposition 1 shows that if there is a minimal critical coalition, then it consists of agents that impose the largest externalities on the others. Next we present a numerical example to show that such a coalition can indeed exists. Consider three airlines and let 1 and 2 be identical. The characteristics of these airlines are such that the only Nash equilibrium is one at which 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 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. We take $q_1 = q_2 = 0.1$ and $q_3 = 0.5$ and $L = 1000$. In addition $p_1 = p_2 = 0.1$ and $c_1 = c_2 = 90$. We do not need to specify p_3 or c_3. In this setting $c_1(\emptyset) = p_1\left[1000 - \frac{0.1x1000}{2} - \left(1 - \frac{0.1}{2}\right)\frac{1000L}{2}\right] = 71.25$. As $c_1 = 90 > c_1(\emptyset) = 71.25$, neither firm 1 nor firm 2 will invest in security if no other firm is investing. And we can clearly pick c_3 big enough that they will not invest either and (N, N, N) is the Nash equilibrium. Now suppose that for some reason outside the model airline 3 changes policy and invests. The returns that airlines 1 and 2 get from investing are now $c_1(\{3\}) = 95$ and as this exceeds the cost of investing both 1 and 2 will change policy if 3 does. Airline three 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 three is the one that imposes the largest externalities on the other airlines in accordance with proposition 1 above.

If we were to tax airlines 1 and 2 enough to cause them to switch policy the tax on not investing would have to be at least $90 - 71.25 = 18.75$. For airline 3 if $p_3 = 0.1$ then we have $c_3(\emptyset) = 90.23$. If $c_3 = 91$ then airlines 3 will not invest if the others do not, but a tax of only $91 - 90.23$ would persuade it to change to investing and

[1] Ties will be broken randomly if needed.

would thus tip the entire equilibrium. The tipping phenomenon thus gives the tax great leverage in this example: it allows taxation to attain the goal of changing the equilibrium at a much lower rate of tax than would otherwise be possible.

The "tipping" issue that we are characterizing here is in fact more general than the particular context in which we are raising it. A general version of the issue that we study is the following. 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 more broadly.

4 Computer Security

The next step is to extend our results to the case of computer security. As noted in our earlier paper the underlying model is the same except that now when a virus affects a computer (the equivalent of a bag with a bomb being loaded by an airline) then it can be transmitted to all other computers on the network and can damage them all rather than damaging only one of them (Anderson [1]). Now we 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** that this is transmitted to all other computers. So clearly $q_i \leq p_i$ and p and q do not refer to independent events. The rest of our notation is the same as before. The only difference from the earlier analysis is that the probability that computer i is infected by computer j is now q_i rather than $q_i/(n-1)$. So for the case of four computers we have

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

and in the general case of n firms when none 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 firms in the set $\{K\}$ are investing in security then the total externality to firm 1 is given by

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

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

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

Here as before $c_1(n,\emptyset)$ is the maximum cost to agent 1 consistent with S being a Nash equilibrium when no other firms invest in security. From our earlier paper we know that if all firms are identical the term $X_1(n,\emptyset)$ goes to L as $n \to \infty$. With heterogeneous agents we cannot establish the same result but it is clear that $X_1(n,\emptyset)$ increases as n increases so 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 5 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.

5 Bankruptcy

There are several more possible applications of the framework that we have developed above. One that we mentioned in the introduction is in the quite different field of corporate finance, and concerns the management of risks in an organization. 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 goes badly wrong, 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. In a rather different line of business Arthur Anderson was recently sent into all-but-bankruptcy by the actions of its Houston branch in managing the Enron audits. In both of these cases we have a situation that 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 large that management closes down this division. So p_i is the risk that a division closes itself down as a result of its risk-taking activities. Let q_i be the probability that division i incurs a loss so large that it is closed down and the entire company is bankrupt and every division is closed. As in the computer security case we must have $q_i \leq p_i$ and p and q do not refer to independent events. The loss in the event of closure of a division is L. Divisions can invest in monitoring their risks at a cost c_i, which is the cost of installing monitoring systems that will detect actions that violate corporate guidelines on risk-taking.[2] Incurring the cost c_i ensures that division will not be closed by its own failures. Clearly when division i takes on a risk it is imposing an external effect on other divisions because there is some chance that they will be closed down because of this. 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 risk is only relevant if the other divisions have not already been closed down by losses originating elsewhere. So we have a problem with the structure of the previous section. As there 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. There is a within-organization Nash equilibrium between the divisions where the strategy choices are whether or not to invest in risk-monitoring. The senior management of the firm would want investing to be a dominant strategy here, but because of the negative externalities it may not be one. They will therefor seek policy measures that will change the payoffs and make investing in risk-control a dominant choice. There will also be the possibility of tipping the equilibrium by persuading a small number of divisions to adopt strict controls.

6 R & D

The same analytical structure can emerge from a radically different problem. Consider a company that needs to solve a problem or to discover new facts. Others are also interested in the same issues - other firms want to solve the same problem or to obtain the same facts. If the problem is solved, or the facts discovered, for one

[2] We could also think of c_i as an opportunity cost resulting from the need to avoid certain deals in order to manage the risk to the enterprise as a whole.

firm, then the solution may be available to some or all others. Information has some of the characteristics of a public good and so may become available beyond the immediate discoverer. In such a situation each firm has to decide whether to invest in obtaining the information or making the discovery. In making this choice it has to bear in mind that another firm might make the discovery and the information might then be available to the choosing firm, rendering its own investment redundant. The investment decision here has the same structure as the problems considered in previous sections.

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 of q_j that another firm j invests and succeeds and that 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 have a payoff of their initial income which as before we take to be Y. If one 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 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

$$p_1 [1 - q_2] G > c_1$$

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 $p_1 G > c_1$. The term $[1 - q_2]$ represents what we previously called contamination and what in this context might also be called the "free rider" effect. It describes the temptation that firm one feels to free ride on firm two's investment in R&D.

There are two possible generalizations of this two by two setup to an n-agent framework: they differ in the assumption made about the dispersal of information once one firm has discovered it. One possibility is that with a certain probability the information is effectively public: the firm might wish to keep it proprietary but the non-rivalrous nature of information dominates and there is a probability $q_j > 0$ that any firm obtains the information from firm j if firm j is successful. If $q_j = p_j$ then the probability that information once discovered becomes public is unity: otherwise for $q_j < p_j$ it is less than one. This is the analog of the assumption that we made in the computer network case that a virus will affect all if it affects one. The alternative is that the information once discovered leaks but leaks to one specific firm only. 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. In what follows we use the former case - there is a chance of information becoming public. In this case each firm undertaking investment provides a positive external effect to the others and in so doing reduces their own incentives to invest. In this case the external benefits to firm one of four firms are as follows:

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

If no other firms invest then the external effects are clearly zero. If firm 2 invests then the expected gain to one is Gq_2 : if two and three invest then we have the expected gain from two plus the expected gain from three conditional on there being no gain from two, etc. This is the same pattern as the computer security case except that we are now dealing with gains rather than losses.

If firms in the set $\{K\}$ are investing in research then the total expected externality to firm 1 is given by

$$X_1(n,\{K\}) = G\sum_{j\notin\{K\}} q_j \prod_{k<j,k\notin\{K\}} (1-q_k) \tag{11}$$

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

$$c_1 < p_1[G-X_1(n,\emptyset)] \equiv c_1(n,\emptyset) \tag{12}$$

Here as before $c_1(n,\emptyset)$ is the maximum cost to agent 1 consistent with I being the best choice for 1 when no 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 agent are identical we do of course have following our earlier paper K-H 2002 that $\lim_{n\to\infty} X_1(n,\emptyset) = G$. From this it follows that as the number of firms increases the incentive that any one has to invest goes to zero.

Within this framework we can now define a Nash equilibrium and enquire into the conditions for tipping as before. The results that emerge are exactly analogous to the computer network case.

7 Conclusions

We have modeled the management of risks that are discrete and interdependent, and have developed a general theory of how groups of firms react to these risks. We have also suggested various policy options for moving equilibria from less to more favorable outcomes. The possibility of tipping introduces novel elements into the policy options. In our earlier review [2] of the homogeneous case we considered options such as insurance, taxation, regulation and standards. All of the comments made there are still relevant to the heterogeneous case, although some refinements are now possible. Consider taxation for example: when all agents are identical then the same tax rate will cause all to switch from not investing to investing in security. Now we have another possibility: a tax may cause members of a minimum critical coalition to switch and thus be instrumental in tipping an equilibrium without itself causing all agents to switch. This may make it possible to reach the same outcome with a lower tax rate than would be needed to cause all agents to switch in the absence of a tipping mechanism. For this to be the case it is necessary that the agents in a minimum critical coalition require lower taxes to induce them to switch than do some other agents. The numerical example after proposition 1 illustrates this case. Similar possibilities arise with respect to the other policy instruments: in each case it may be sufficient to introduce the instrument at a level that will affect the behavior of a critical coalition.

In addition the tipping possibility adds a separate dimension into the policy options. It allows us to identify a set of agents whose behavior is strategic. To change the equilibrium it is sufficient to change the behavior of these agents as all others will follow. Changing the behavior of a small group is a qualitatively different task from changing that of a large group so this alters the policy landscape substantially.

References

[1] 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

[2] Kunreuther Howard and Geoffrey Heal. Interdependent Security: the Case of Identical Agents. Working Paper, Columbia Business School.

[3] Schelling, Thomas. (1978) Micromotives and Macrobehavior New York: Norton