Interdependent Security: A General Model

Geoffrey Heal and Howard Kunreuther[*†]
Columbia Business School
Wharton School, University of Pennsylvania

July 25 2004

Abstract

In an interdependent world the risks faced by any one agent depend not only on its choices but also on those of all others. Expectations about others' choices will influence investments in risk-management, and the outcome can be sub-optimal investment all round. We model this as the Nash equilibrium of a game and give conditions for such a sub-optimal equilibrium to be tipped to an optimal one. We also characterize the smallest coalition to tip an equilibrium, the minimum critical coalition, and show that this is also the cheapest critical coalition, so that there is no less expensive way to move the system from the sub-optimal to the optimal equilibrium. We illustrate these results by reference to airline security, the control of infectious diseases via vaccination and investment in research and development.

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

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

---

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

[†]We are grateful to Douglas Bernheim, Charles Calomiris, Avinash Dixit, Michael Kearns, Nat Keohane, Paul Kleindorfer, Erwann Michel-Kerjan, Peter Orszag, Todd Sandler, Joe Stiglitz, Thomas Weber, Richard Zeckhauser anf two referees for valuable comments on various drafts of this paper. We have also benefitted from the comments of participants in NBER Insurance Workshops, the AEA annual meetings in 2003, and seminars at Columbia, Stanford, University of Pennsylvania, Cornell, Ohio State, the London School of Economics and the U.S. General Accounting Office. Responsibility for errors is of course our own. We acknowledge financial support from Radiant Trust, Lockheed Martin, the Wharton Risk Management and Decision Processes Center and the Columbia Earth Institute.

1 Introduction

The problem structure that we study was motivated by examining the risks associated with terrorism, though as we shall indicate the concept of interdependent risks that emerged from this analysis is a very general one. The central issue is behavior in the face of risks whose magnitude depends on an agent's own risk-management strategies and on those of others. We have called this class of problems **interdependent security (IDS) problems** and have used game-theoretic models to characterize their Nash equilibria [13] and [9] .

Risks of terrorism are typically interdependent as the risks to which one organization is exposed depend not only on its own choice of security investments, but also on the actions of other agents. Failures of a weak link in an interdependent system can have devastating impacts on all parts of the system. Because interdependence does not require proximity, the antecedents to catastrophes can be quite distinct and distant from the actual disaster, as in the case of the 9/11/01 attacks, when security failures at Boston's Logan airport led to crashes at the World Trade Center (WTC), the Pentagon, and in rural Pennsylvania. The same was true in the case of the August 2003 power failures in the northeastern US and Canada, where the initiating event occurred in Ohio, but the worst consequences were felt hundreds of miles away. Similarly a disease in one region can readily spread to other areas with which it has contact, as was the case with the rapid spread of SARS from China to its trading partners.

Investing in airline security is a clear example of an IDS problem. Even the adoption of elaborate security procedures by one air carrier may not mitigate the risks faced due to the baggage or passenger transfers from other less diligent airlines. Under some conditions, the added risk from others' lax inspections reduces the benefits to diligent airlines from their strict inspections to the point where the costs of such inspection can no longer be justified by the expected benefits. In equilibrium, all actors may fail to invest in strict security measures.

Lest this point be taken as theoretical, recall the explosion of Pan Am 103 over Lockerbie, Scotland, in December 1988. In Malta terrorists checked a bag containing a bomb on Malta Airlines, which had minimal security procedures. The bag was transferred at Frankfurt to a Pan Am feeder line and then loaded onto Pan Am 103 in London's Heathrow Airport. The transferred piece of luggage was not inspected at either Frankfurt or London, the assumption in each airport being that it was inspected at the point of origin. The bomb was designed to explode above 28,000 feet, a height normally first attained on this route over the Atlantic Ocean. Failures in a peripheral part of the airline network, Malta, compromised the security of a flight leaving from a core hub, London.

Not only do interdependencies mean that one agent's exposure depends on the actions of others, but this interaction can reduce the incentive that any agent has invest. If an airline receives bags and passengers from other airlines that do not check them thoroughly, the expected benefits from its own investment in security is compromised, as these uninspected bags mingle with their own inspected bags.[1] (Note that we assume that an airline does not check the bags transferred from other airlines due to time and constraints. Until recently this was true for all airlines except El Al.) If the incentives to invest are thus reduced all round, the outcome may be a situation with dramatically sub-optimal investment in security.

2 Classes of IDS Problems

Interdependent security (IDS) problems have one common characteristic—the decision by one agent on whether or not to incur an investment cost will impact on the welfare of other agents, and will also affect their incentives to invest in prevention. Next we categorize three classes of IDS problems based on differences in their Nash equilibria.

2.1 Class 1: Partial Protection

An agent's investment to reduce its own risks also decrease the risks experienced by others. The more agents that invest in preventive measures, the lower are the negative externalities in the system. To take the example of airline security discussed above, if airlines face terrorist risks and Airline 1 invests in a stricter baggage screening system, then all the other airlines benefit because they now have a smaller chance of receiving a transferred bag that contains a bomb. The more airlines that increase their investment in baggage security, the greater the reduction in the risk experienced by everyone else in the system.

A situation where an agent knows that there is a chance that others will still subject it to risk even if it invests in protection is a Class 1 problem. For example, an apartment owner considering investing in fire prevention equipment has to take into account the possibility that a fire from a nearby unprotected apartment will spread to her unit even if she invests in risk-reducing measures. As the number of apartments investing in fire prevention equipment increases, the likelihood that her apartment will suffer a fire loss from others decreases. She will then have more of an economic incentive to incur these investment costs herself. The decision by electric utilities to invest in measures to reduce the likelihood of a power failure is

[1] We talk here mainly of inspecting bags but the same issue arises with passengers - indeed according to terrorism expert Gordon Woo of Risk Management Solutions, "As of a few months ago, all transatlantic bookings are screened for suspicious passengers. For example, when making a BA booking, provision of passport information is now mandatory. The problem is with transit passengers, whose security status may be inadequately assessed at their original point of departure. Unlike baggage, it can take many hours to perform a passenger security check. When BA223 restarted after the New Year cancellations, there were delays of hours as checks were made on passengers. Delaying a flight by several hours is intolerable. The current policy is thus to cancel flights rather than delay them." Personal communication, February 16 2004.

also partially determined by what others do. Each utility knows that there is some chance that an outage in another part of the country can knock out its power even if it has undertaken its own preventive measures. (Heal [10]).

2.1.1 Nash Equilibria

As we show more formally in the next section, this class of problems can have multiple Nash equilibria. The most interesting case is where there are two equilibria: either all agents invest in security or none of the agents want to do this. Then there is the possibility of tipping or cascading: inducing some agents to invest in prevention will lead others to follow suit.

2.1.2 Private and Social Welfare

Whenever there are two Nash equilibria involving everyone or no one investing in security, then the socially optimal solution will always be for everyone to invest. Each agent will find that the cost of investing in protection will be justified if it doesn't incur any negative externalities and society will be better off as well.

When there is only a single Nash equilibrium, there are some situations where the investment choices by agents are efficient. The most obvious one is where the costs of protection are sufficiently low so that each agent wants to invest in protection even when all the other agents had decided not to incur these costs. If the costs of investment to each agent are very high, then it may be efficient for no one to incur them; however, there are cases when the costs are high enough that each agent does not want to invest in protection, but it would be better for society if some or all of them did so. A formal treatment of these and other cases appears in Sections $3-5$, and a set of illustrative examples with respect to airline security are presented in Section 6.

2.2 Class 2: Complete Protection

This class of IDS problems differs from Class 1 in that if an agent invests in security then it cannot be harmed at all by the actions or inactions of others and reciprocally it cannot affect others. As an example, a completely effective vaccine will protect a person against catching a disease from contagious individuals. Prior to getting vaccinated this person may be susceptible to the disease and could infect others. An airline, such as El Al, that checks both its own bags and those transferred from other airlines also illustrates Class 2 behavior.

A related example from the field of organizational decision-making is where a division in a firm decides to incur the cost of separating itself from the rest of the organization (e.g. as a captive) so that it cannot be hurt by other divisions and cannot harm them if it suffers a loss. As more divisions decide to take such action, then this decreases the likelihood of any unit that is still part of the larger organization being economically harmed by others. Breaking away in this manner then becomes less attractive.

2.2.1 Nash Equilibrium

For Class 2 problems there is only one Nash equilibrium and it can range from the extremes of either all agents or no agents adopting security, with intermediate cases where some agents invest. Since there is only one Nash equilibrium it is impossible to have tipping or cascading in Class 2 problems. In fact, it is less attractive for an agent to invest in protection, should others then decide to do so. In deciding whether or not to invest, Agent i compares the expected benefits and costs. As more agents invest, the expected benefit to i of following suit decreases, since there is a reduction in the negative externalities which translates into a lower probability of suffering a loss.

2.2.2 Private and Social Welfare

As in Class I problems the number of agents investing at a Nash equilibrium will not exceed the number that would be socially optimal. Each agent i does not take into account the negative externalities it is creating in determining whether to invest in protection. For the situation where the investment costs are so low that every agent will want to protect itself, then the Nash equilibrium will be efficient for the same reasons as it is for Class 1 problems. Similarly, one could have an efficient Nash equilibrium where no one invests in protection because the costs of taking this action are so high. On the other hand, there can be a range of parameters where the Nash equilibrium will not be socially optimal. We will discuss the vaccination problem in Section 7 of the paper.

2.3 Class 3: Positive Externalities

For this class of problems an investment by one agent creates positive externalities, making it less attractive for others to follow suit. A firm's decision on whether to incur expenditures for research and development (R&D) will be partially influenced by what other firms in the industry are doing. Suppose firm i has decided to invest in R&D and firm j has to decide whether to do likewise. The greater the likelihood that j can benefit from the success of i, the less likely it is that j will invest in R&D. Class 3 problems include situations where there is investment in knowledge and agents can learn from successful investments by others.

2.3.1 Nash Equilibrium

As is the case for Class 2 problems, there is only a single Nash equilibrium here, but for a very different reason. As more agents invest in knowledge, there is a greater chance that those on the sidelines will be able to benefit from their successes (i.e. there is an increase in positive externalities). For this class of problems you cannot have tipping and cascading: if any agent convinces others to invest, it will have less rather than more reason to do so itself.

2.3.2 Private and Social Welfare

A Nash equilibrium is efficient if the only agents who do not invest are those for whom the expected benefits to themselves and others do not exceed the cost of the investment. There will be situations where the costs of investment is sufficiently high that an agent will not want to incur it even though by doing so other agents in the system will benefit. In this situation there will be fewer agents in equilibrium investing in knowledge than would be socially optimal. We will discuss the R&D problem in Section 8 of the paper.

3 The Model

In this paper we present a general model of IDS problems, which covers all three classes of problems discussed above. We characterize Nash equilibria, show that they exist, specify conditions for the existence of multiple equilibria, one of which involves investment in security by all agents while the other involves no investment by any agents. We then characterize the possibility of tipping and cascading the equilibria from a state of no investment to one of universal investment in security.[2] We define a critical coalition as one where a change from not investing to investing by its members will induce all non-members to follow suit. We then characterize the properties of minimum critical coalitions in terms of Pigouvian externalities, show that it is generically unique and identical to the (unique) cheapest critical coalition. Strategic complementarity and substitutability (Bulow Geanakoplos and Klemperer [1]) lie at the heart of some of the phenomena that we study.

We consider A interdependent risk neutral agents indexed by i. Each is characterized by parameters p_i, L_i, c_i and Y_i. Here p_i is the probability that agent $i's$ actions lead to a direct loss L_i. A direct loss can be avoided with certainty by investing in loss-prevention at a cost of c_i. Initial income before any losses are incurred or before expenditure on loss-prevention is Y_i. Each agent i has a discrete strategy, X_i, that takes as values either S or N representing investing and not investing respectively. If i incurs a direct loss, then this may also affect other agents' outcomes. We call the loss (or in some cases gain) to them in this case an indirect impact. More specifically $q_i\left(K, X_j\right)$ is the expected loss to agent i when it follows strategy X_i and the agents in the set $\{K\}$ invest in loss-prevention, with those not in $\{K\}$ not investing. When we use a letter to refer to a set we will designate it $\{K\}$, except when it is an argument

[2] For discussions of tipping and cascading in the literature, see Schelling [16], Dixit, [5], Watts [17] and Gladwell [6].

of a function, in which case we omit the parentheses. A feature of the IDS problem described above is that an agent who has invested in prevention cannot cause an indirect impact on others, so if $\{K\} = \{1, 2, ..i-1, i+1, ..A\}$ then $q_i(K, X_i) = 0$ whether $X_i = S$ or N.

If agent i invests in prevention and agents in the set $\{K\}$ are also investing then the expected cost from this is $c_i + q_i(K, S)$ where the first term is the direct cost of investing and the second is the expected cost (or benefit if negative) of indirect impacts imposed by others who do not invest. The expected cost of not investing is given by $p_i L + (1 - \alpha p_i) q_i(K, N)$. Here the first term is just the expected direct loss and the second is the expected indirect impact. In this second term the parameter $\alpha \in [0, 1]$ indicates the extent to which damages are non-additive. If $\alpha = 0$ then this second term is $p_i L_i + q_i(K, N)$, so that the total expected damage sustained by agent i in the case of non-investment is the sum of the direct and indirect effects. If however $\alpha = 1$ then we have $p_i L + (1 - p_i) q_i(K, N)$ which means that the indirect effects are conditioned on the direct loss not occurring. In this case the damages from harmful events are non-additive (i.e., you only die once). A second plane crashing into one of the towers of the World Trade Center would not have increased the damage from 9/11 significantly, and a second bomb placed on PanAm 103 would likewise have inflicted no extra damage.

The agent is indifferent between investing and not investing when

$$c_i + q_i(K, S) = p_i L_i + (1 - \alpha p_i) q_i(K, N) \tag{1}$$

or

$$c_i(K) = p_i L_i + (1 - \alpha p_i) q_i(K, N) - q_i(K, S) \tag{2}$$

The value of the cost given by equation (2), $c_i(K)$, is the cost of investment at which i is just indifferent between investing and not investing: if $c_i < c_i(K)$ then she will invest and vice versa.

The IDS problems associated with airline security that we first studied were Class 1 problems [9] [13] where

- $q_i(K, N) = q_i(K, S)$ and $\alpha = 1$

so that

$$c_i(K) = p_i(L_i - q_i(K, N)) \tag{3}$$

It follows in this case that $c_i(K)$ increases in K: as more agents invest then the expected indirect loss falls and the cost threshold for investment rises, with $c_i(\emptyset) < c_i(A - i)$, the latter being the critical cost when all agents other than i are investing. In Class 2 problems where there is complete protection, such as deciding whether to get vaccinated, then

- $q_i(K, S) = 0$ whatever the set $\{K\}$ and $\alpha = 1$

so that

$$c_i(K) = q_i(K, N)[1 - p_i] + p_i L_i \tag{4}$$

Here $c_i(K)$ decreases with K, so that $c_i(\emptyset) > c_i(A - i)$. This reflects an important difference between these two cases, which is the sign with which q_i enters on the RHS of the equation, negative in one case and positive in the other. In Class 3 problems, such as determining whether to invest in R&D,

- both L_i and q_i are positive, $q_i(K, N) = q_i(K, S)$ and agents invest to generate benefits to themselves and others.

Those who invest in R&D may make a discovery. If they do, then they provide spillover benefits to others who may not have invested in R&D, so that

$$c_i(K) = p_i[q_i(K, S) - L_i] \tag{5}$$

Here again $c_i(K)$ decreases with K.

We now investigate properties of the Nash equilibria of this system.

Definition 1 *A Nash equilibrium is a set of strategies* $X_1, ...X_A$ *such that (1)* $X_i = S$ *for all* $i \in \{K\}$ *(which may be empty), (2) if* $X_i = S$ *then* $c_i(K) \geq c_i$ *and (3) if* $X_i = N$ *then* $c_i(K) \leq c_i$.

Theorem 2 *A Nash equilibrium in pure strategies exists.*

Proof. We prove existence of an equilibrium constructively,[3] giving an algorithm which will terminate by locating an equilibrium.

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

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

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

[3] This argument is based on a discussion with Michael Kearns, personal communication [11].

There may be equilibria where all agents invest in loss-prevention, those where none do, and mixed equilibria where some invest and others do not. We will illustrate these equilibria in the context of the airline security example in Section 6. It is also possible that for some parameter values there is more than one equilibrium as the following proposition indicates:

Proposition 3 *There are Nash equilibria at which all agents invest and also Nash equilibria at which none invest if and only if* $c_i(\emptyset) < c_i < c_i(A-i)\,\forall i$. *If both* $(N, N, .., N)$ *and* $(S, S, ..., S)$ *are Nash equilibria, then* $(S, S, , .., S)$ *Pareto dominates* $(N, N, ..., N)$.

Proof. First note that $c_i(A-i) > c_i(\emptyset)$ for the Class 1 IDS problems, so the conditions of the proposition are not vacuous. If $c_i > c_i(\emptyset)$ then $X_i = N\forall i$ is an equilibrium because it satisfies the definition with $\{K\} = \emptyset$. And if $c_i < c_i(A-i)$ then $X_i = S\forall i$ is an equilibrium with $\{K\} = \{A\}$. Conversely if $X_i = N\forall i$ is an equilibrium then $c_i > c_i(\emptyset)\,\forall i$ and if $X_i = S\forall i$ is an equilibrium then $c_i < c_i(A-i)\,\forall i$. This proves the first part of the proposition.

The proof of the second part is as follows. From equation (1), Pareto domination by the $(S, S, .., S)$ equilibrium is equivalent to

$$c_i < p_i L_i + (1 - \alpha p_i) q_i(\emptyset, N) \tag{6}$$

where the LHS of (6) reflects the costs to each agent i if all agents invest in prevention and the RHS of (6) is the cost to agent i if no-one invests in prevention. The existence of both $(S, S, .., S)$ and $(N, N, ...N)$ as equilibria implies that

$$c_i < c_i(A-i) = p_i L_i + (1 - \alpha p_i) q_i(A-i, N)] - q_i(A-i, S)] \tag{7}$$

The RHS of (7) is less than the RHS of (6) because $q_i(\emptyset, N) > q_i(A-i, N)$ and $q_i(A-i, S) \geq 0$ so that (7) implies (6). This completes the proof of the proposition. ■

If there are two equilibria, one with all not investing and the other with everyone investing in protection, then it is obviously interesting to know how we might tip the inefficient $(N, N.....N)$ equilibrium to an efficient $(S, S, ...S)$ equilibrium. Next we look into the possibility of tipping the non-investment equilibrium.

3.1 Tipping

Definition 4 *Let* $X_i = N\forall i$ *be a Nash equilibrium. A critical coalition for this equilibrium is a set* $\{M\}$ *of agents such that if* $X_i = S\forall i \in \{M\}$ *then* $c_j(M) \geq c_j \forall j \notin \{M\}$. *A minimum critical coalition is a critical coalition with two additional properties: no subset is a critical coalition, and no other critical coalition contains fewer agents.*[4]

[4]See Heal [8] for an earlier use of the idea of a minimum critical coalition in a different context, unconnected with tipping.

Define

$$q_i^j(K,N) = q_i(K-j,N) - q_i(K,N) \geq 0 \tag{8}$$

This is the change in the expected indirect loss to agent i, who does not invest in loss-prevention, when agent j joins the set $\{K\}$ of agents who are already investing in loss-prevention. For the remainder of this section we make the following assumption:

*Assumption A*1: $q_i^j(K,N)$ is independent of i: $q_i^j(K,N) = q^j(K,N)\,\forall i$

This implies that indirect effects are symmetrically distributed across agents. Also define $q_i^j(\emptyset,S) = q_i(\emptyset,S) - q_i(j,S)$ and $q_i^j(\emptyset,N) = q_i(\emptyset,N) - q_i(j,N)$ and make the additional assumption that

$$\textit{Assumption A2: } q_i^j(\emptyset,S) = q_i^j(\emptyset,N) = q_i^j(\emptyset) = q^j(\emptyset) \tag{9}$$

This indicates that the indirect impact of a change of strategy by agent j on another agent does not depend on the other agent's strategy.

Finally, we shall need the following assumption:

$$\textit{Assumption A3: } \text{The ranking of agents by } q^j(K) \text{ is independent of } \{K\} \tag{10}$$

This says in intuitive terms that if agent k creates the largest negative externalities when agents in the set $\{K\}$ are investing in loss-prevention, then agent k creates more externalities than any other agent whatever the set investing in loss prevention.

Theorem 5 *Let $X_i = N\forall i$ be a Nash equilibrium. If a minimum critical coalition exists for this equilibrium then for some integer K it consist of the first K agents when agents are ranked in decreasing order of $q^j(\emptyset)$.*

Proof. Recall from (2) that $c_i(K) = \{q_i(K,N) - q_i(K,S)\} + p_i(L_i - \alpha q_i(K,N))$ and define

$$\Delta c_i^j = \{q_i(j,N) - q_i(j,S)\} + p_i(L_i - \alpha q_i(j,N)) - \tag{11}$$

$$\{q_i(\emptyset,N) - q_i(\emptyset,S)\} - p_i(L_i - \alpha q_i(\emptyset,N)) \tag{12}$$

$$= (1-\alpha)\{q_i(j,N) - q_i(\emptyset,N)\} + \{q_i(\emptyset,S) - q_i(j,S)\} \tag{13}$$

Using $A3$, we see that for $\{K\} = \{1,2,3,...k\}$ to form a critical coalition (where agents are ranked in decreasing order of $q^j(\emptyset)$) it must be the case that k is the first integer such that

$$\sum_{j=1}^{j=k} \Delta c_i^j \geq c_i - c_i(\emptyset)\,\forall i > k \tag{14}$$

which can be written as

$$\sum_{j=1}^{j=k} \left(q^j(\emptyset,S) - (1-\alpha)\,q^j(\emptyset,N)\right) \geq c_i - c_i(\emptyset)\,\forall i > k \tag{15}$$

By (9) this can be simplified to

$$\alpha \sum_{j=1}^{j=k} q^j(\emptyset) \geq c_i - c_i(\emptyset)\, \forall i > k \tag{16}$$

This completes the proof. ■

Corollary 6 *There is a minimum critical coalition only if $\alpha > 0$, i.e., if there is some degree of non-additivity of damages.*

Corollary 7 *A minimum critical coalition is unique if the values of the quantities $q^j(K)$ are different for different agents j.*

Proof. This follows immediately: the only way a MCC could not be unique is if $q^j(K) = q^f(K)$ for some $j, f < K$, which is ruled out by the assumption. ■

These results imply that a MCC is easily characterized and that in general there is only one MCC, as generically with respect to parameter values the terms $q^j(K)$ will differ. Note that assumption (9) simplifies the formula but is not necessary for a result of this type: without it we would have to rank agents by $(q^j(S) - (1-\alpha)\, q^j(N))$, which simplifies to $q^j(S)$ if $\alpha = 1$ and to $q^j(S) - q^j(N)$ if $\alpha = 0$.

4 Minimal, Cheapest & Pigouvian Critical Coalitions

What are the policy implications of our results on minimal critical coalitions? Clearly one is that an equilibrium with no investment in security may be converted to one with full investment by persuading a subset of the agents to change their policies. But is this the best way to change equilibrium? Would other alternatives be less costly? To investigate this we need to calculate the cost of incentives sufficient to induce agents to switch strategy from N to S at an equilibrium where all are choosing N.

There are a number of ways in which one can approach this issue. We assume that, as provided for in Proposition 3, there are (at least) two Nash equilibria, one with $X_i = N \forall i$ and the other the opposite: $X_i = S \forall i$. The objective of policy is to support the equilibrium where all agents choose S, because this Pareto dominates. The regulatory authority or an industry trade association is therefore interested in knowing the least expensive way of changing an equilibrium with no investment into one where all invest. It is clear that insights into the minimum coalitions that will tip the others plays a role in the analysis.

Tipping assures us that rather than providing incentives to all agents, we can work with a subset and that it is only necessary to persuade them to alter their behavior. The least expensive way of changing equilibrium is therefore likely to involve providing incentives for a critical coalition to change its behavior and tip the entire system. Will it involve working with a minimum critical coalition? To address this question one needs to understand the relationship between an MCC and the cheapest critical coalition (CCC).

If $(N, N, ..., N)$ and $(S, S, ..., S)$ are two Nash equilibria, how do we compute the cost to the regulator of persuading one agent, say agent i, at the $(N, N, ...N)$ equilibrium to change its strategy to S? The answer depends on what assumption i makes about the behavior of the other agents. We consider the case in which i assumes that it is the only agent to change strategy, so that it will assess the cost as

$$Y_i - p_i L_i - (1 - \alpha p_i) q_i(\emptyset, N) - [Y_i - c_i - q_i(\emptyset, N)] = c_i - c_i(\emptyset) > 0 \qquad (17)$$

The cost of persuading a critical coalition C to change strategy is $\sum_C [c_i - c_i(\emptyset)]$.

Recall that from our earlier results a minimum critical coalition of k agents is a group consisting of the first k agents ranked by the index $q^j(\emptyset)$. We assume again that this index is independent of the agent i and of that agent's strategy. We now show that a minimum critical coalition also a cheapest critical coalition. *A priori* this is not obviously the case since there is no intuitive connection between the policy problem (17) and the criterion for selecting a MCC, which is based on reducing the negative externalities in the system (i.e., $\sum_{j \in M} q^j(\emptyset)$) using as few agents in the set$\{M\}$ as possible. More specifically we can prove the following result:

Theorem 8 *Any cheapest critical coalition (CCC) is also a minimal critical coalition (MCC): equivalently the set of CCCs is contained in the set of MCCs.*

Proof. The proof is by contradiction. Let $\{C\}$ be a cheapest critical coalition and assume contrary to the theorem that it is not a MCC. Then there exists at least one agent say agent l who can be removed from the coalition without it losing criticality. But this reduces the cost as the sum in (17) now excludes l and is over one less agent, and each component of the sum is positive. Hence the coalition $\{C\}$ could not have been the cheapest, a contradiction. ■

Recall that provided that the terms $q^j(\emptyset)$ differ from agent to agent, a MCC is unique: *in this general case we know that there is a unique CCC and that it is equal to the MCC.*

This result establishes an unexpected connection between two lines of argument - one concerning the smallest critical or tipping coalition in a numerical sense and the other concerning the least expensive critical coalition in the sense of what it costs to induce this coalition to change its policies from not investing to investing in loss-prevention. The parameters that make for membership of an MCC do not address the cost of inducing a strategy change and so one does not naturally expect a connection, yet there is one: all CCCs are also MCCs, and indeed in general the unique MCC is also the unique CCC.

Finally we compare the tipping costs of moving from an equilibrium with no investment to one with full investment via the CCC with the cost of attaining the same outcome via Pigouvian subsidies - a standard approach within the public finance tradition given the positive external effects associated with investment in loss-prevention. The difference between the private and social benefits of investment by agent j when no other agents are investing is given by the negative externalities created by j [i.e., $q^j(\emptyset)$]. This is therefore the Pigouvian subsidy that should be offered to agent j to persuade it to invest at an equilibrium where no one else is investing. The Pigouvian subsidy and the size of the negative externality used to rank members of an MCC are identical. Identifying the members of the MCC has value to the government or a trade association that wants to maximize industry profits. The reduced subsidy to achieve full investment is reflected in the difference between the cost of subsidizing all agents and that of subsidizing just the members of the MCC.

5 Endogenous Probabilities

In the above analyses the risks faced by the agents are assumed to be independent of their behavior. In reality if some agents are known to be more security-conscious than others, they are presumably less likely to be terrorist targets. There is a resemblance here to the problem of theft protection: if a house announces that it has installed an alarm, then burglars are likely to turn to other houses as targets [13]. In the case of airline security, terrorists are more likely to focus on targets which are less well protected, so that p_i depends on whether or not agent i invests in security. This is the phenomenon of displacement or substitution, documented in Sandler [15]. Keohane and Zeckhauser [12] also consider the implication of endogenous terrorist risks, focusing on ways to controlling the stock of terror capital and curbing the flow into the terrorist organizations.

We focus here on the case of airline security and assume that the risk faced by an airline that does not invest in stricter inspections increases as the fraction of airlines investing in such measures increases. In other words, if more airlines from a given population invest in security, then those who do not take similar actions become more vulnerable. Formally let $\eta(S)$ be the number of airlines investing in security, i.e. the number in the set $\{K\}$ of airlines that are investing. The relevant probabilities facing those firms not investing in security, $p_i(\eta(S))$, are now increasing in $\eta(S)$.

Now return to equation (1) 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 and p_i are exogenous. Assume as before that $q_i(K,N) = q_i(K,S)$.

$$c_i(K) = p_i(\eta(S))\left[L_i - \alpha q_i(K,N)\right] \tag{18}$$

In this expression $q_i(K,N)$ depends on $\eta(S)$ since the likelihood of airline i being impacted by others depends how many airlines are investing in security. To understand how a change in $\eta(S)$ will affect $c_i(K)$, assume that K is large enough to be treated as a continuous variable and differentiate the right hand side of (18) with respect to $\eta(S)$:

$$\frac{dc_i(K)}{d\eta} = [L_i - \alpha q_i] \frac{dp_i}{d\eta} - \alpha p_i \frac{dq_i}{d\eta} \tag{19}$$

Here $\frac{dp_i}{d\eta} > 0$ by assumption, and the coefficient on this is the difference between the direct and expected indirect losses, which we assume to be positive. The first term on the RHS of (19) is therefore positive under this assumption. The term $\frac{dq_i}{d\eta}$ measures the impact of a change in the number investing on the total expected indirect impact on firm i, a non-investor. We assume this to be negative: more agents investing in security means less exposure to indirect effects. This certainly seems reasonable for the airline case.

Theorem 5 on tipping is relevant to a model with endogenous probabilities. With the above assumptions, $\frac{dc_i(K)}{d\eta} > 0$ in (19) and an increase in the number of agents investing in security will raise the threshold cost level for the remaining agents to invest, thus making it more likely that they will also invest. 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 increases the risk that a firm that did not invest in security will become a terrorist target.

Theorem 2 which shows that there exists a Nash equilibrium in pure strategies for the case of exogenous probabilities also holds for the case of endogenous probabilities given the above assumptions. For the argument to work we require that it still be the case that a firm is most likely to choose S when all others are also choosing S and that if in such a situation it chooses N then it will always choose N. But this is implied by the assumption that the total externality imposed on a firm decreases as the number of other firms investing increases.

6 Class 1 Illustrative Examples: Airline Security

This section provides an analysis of Nash equilibria for Class 1 problems by focusing on airline security. We analyze the relationship between private and social welfare and illustrate tipping and cascading through numerical examples coupled with a geometric framework to provide intuition for these results. There are 2 separate airlines. To simplify the computations we utilize a model that is less general than that of section 3 and assume that the probabilities of a terrorist attack are known and fixed. Let p_{ij} be the probability that on any trip a bag containing a bomb is loaded onto airline i and is then transferred to airline j and explodes on j. If $i = j$, we have the probability that an airline loads a bag with a bomb and this explodes on its own plane. Each airline can either invest in a security system S at a cost per trip of $c_i > 0$ or not invest N. Security systems are assumed to be completely effective so that they eliminate the chance of a bomb coming through the airline's own facility. In the event that a bomb explodes on a plane the loss is $L > 0$. The initial income of an airline is $Y > c_i \; \forall i$.

This framework gives rise to the following payoff matrix showing the outcomes for the four possible combinations of N and S. If both airlines invest in security systems then their payoffs per trip are just their initial incomes net of the investment costs. If A_1 invests and A_2 does not, then A_1 has a payoff of income Y minus investment cost c_1 minus the expected loss from a bomb transferred from A_2 that explodes on A_1 (i.e., $p_{21}L$), while A_2 has a payoff of income Y minus the expected loss from a bomb loaded and exploding on its plane, $p_{22}L$. If neither invests then A_1 has a payoff of income Y minus the expected loss from a bomb loaded and exploding on its own plane $p_{11}L$ minus the expected loss from a bomb transferred from A_2 that explodes on A_1 (i.e., $p_{21}L$) conditioned on there being no explosion from a bomb loaded by A_1 itself $(1 - p_{11})$. A_2's payoff is determined in a similar fashion.

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

Within this framework assumptions A1 to A3 are always satisfied, as is the condition $c_i(\emptyset) < c_i < c_i(A)\, \forall i$ of proposition 1. We also have $q_i(K, N) = q_i(K, S)$ and $\alpha = 1$.

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

$$c_1 < p_{11}L \text{ and } c_1 < p_{11}[1 - p_{21}]\, L \tag{20}$$

The condition that $c_1 < p_{11}L$ is clearly what we would expect from a single airline operating on its own. The tighter condition that $c_1 < p_{11}[1 - p_{21}]\, L$ reflects the risk imposed by a firm without security on its competitor: this is the risk that dangerous baggage will be transferred from an unsecured airline to the other.

The nature of the Nash equilibrium in the interdependent security model naturally depends on the parameters. From the payoff matrix it is clear that (S, S) is a Nash equilibrium if $c_i < p_{ii}L$ and is a dominant strategy if $c_i < p_{ii}L(1 - p_{ji})$ where i and j are 1 or 2. (N, N) is a Nash equilibrium if $c_i > p_{ii}L(1 - p_{ji})$ and a dominant strategy if $c_i > p_{ii}L$. From these inequalities we note that (S, S) and (N, N) are both Nash equilibria if $p_{ii}L(1 - p_{ji}) < c_i < p_{ii}L$: this is consistent with proposition 3, as $p_{ii}L(1 - p_{ji}) = c_i(\emptyset)$ and $p_{ii}L = c_i(A - i)$. Finally if $c_1 > p_{11}L$ but $c_2 < p_{22}L(1 - p_{12})$ then (N, S) is a Nash equilibrium, and if 1 and 2 are interchanged then the equilibrium is (S, N). This configuration of Nash equilibria is summarized in Figure 1. Note that if $c_1 = c_2$ then we are on the diagonal of figure 1 and the only possible equilibria are (S, S), either (S, S) or (N, N), and (N, N). In this case mixed equilibria are not possible, as stated in our earlier paper [13].

Figure 1 enables one to determine when the Nash equilibrium also is socially optimal. From the proof of Proposition 3 we know that when both (S, S) and (N, N)

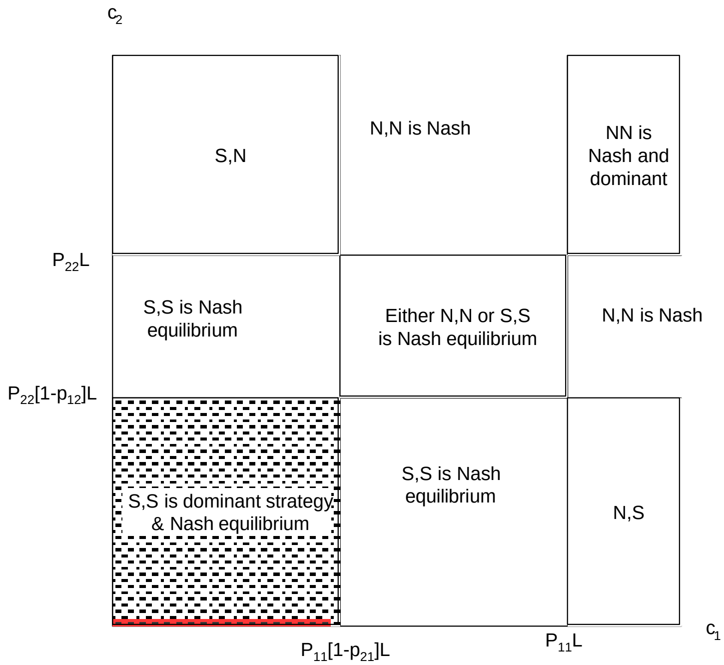

Figure 1: Nash equilibria as a function of c_1 and c_2.

equilibria coexist the former Pareto dominates the latter, and to be precise we know that (S,S) Pareto dominates (N,N) whenever

$$c_i < p_i L_i + q_i(\emptyset, N) - \alpha p_i q_i(\emptyset, N) \tag{21}$$

In the present context this simplifies to

$$c_i < p_{ii} L + L(1 - p_{ii}) p_{ji} \tag{22}$$

In terms of figure 1, this means that the area in which (S,S) Pareto dominates is a rectangle that includes but is greater than the region that is the product of the two intervals $[0, p_{11}L]$ and $[0, p_{22}L]$. So this includes all of the regions in which (S,S) is a Nash equilibrium, and parts of the regions in which (S,N), (N,S) and (N,N) are equilibria. In particular, whenever (S,S) is an equilibrium, then it is efficient.

6.1 Tipping

Consider three airlines, and let $p_{11} = p_{12} = p_{13} = p_{21} = p_{22} = p_{23} = 0.1$, $p_{31} = p_{32} = 0.3, p_{33} = 0.2$, $L = 1000$ and $c_1 = c_2 = 85, c_3 = 200$. The Nash equilibria for this problem are depicted in figures 2 and 3. In this setting

$$c_1(\emptyset) = c_2(\emptyset) = p_{11} L (1 - p_{21} - (1 - p_{21}) p_{31}) = 63$$

As $c_1 = c_2 = 85 > c_1(\emptyset) = c_2(\emptyset) = 63$, neither firm 1 nor firm 2 will invest in security if firm 3 is not investing. We have that

$$c_3(\emptyset) = p_{33}L(1 - p_{23} - (1 - p_{23})\,p_{13}) = 162$$

and so firm three will not invest (as $c_3 = 200$) and (N, N, N) is the Nash equilibrium. If firm 3 does not invest, then not investing is a dominant strategy for both the other firms for any cost above 63.

Suppose that airline 3 is required to invest in security by either an airline association or the federal government. It now imposes no externality on the other firms and so does not affect their decisions. To understand the choices of firms 1 and 2 we simply have to apply inequality (20), which gives a critical cost level of 90, meaning that investment will now be a dominant strategy when the cost is less than 90. As the actual cost for firms 1 and 2 is less than this by assumption at $c_1 = c_2 = 85$, we see that after firm 3 has changed strategy from N to S the dominant strategy for both firms 1 and 2 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.

The tipping phenomenon is shown geometrically in figures 2 and 3. These are similar to figure 1 above, showing the sets of $\{c_1, c_2\}$ values corresponding to different equilibrium types. The key point in seeing tipping geometrically is that this diagram for firms 1 and 2 depends on what firm 3 does. A change by 3 alters the entire equilibrium diagram for the other two firms.[5] 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.

We now compute the expected profits of each airline when none of the airlines are investing in security. The expected loss for airline 1 (and also for 2) at an equilibrium where no firms invest is

$$p_{11}L + (1 - p_{11})\,[p_{21}L + (1 - p_{21})\,p_{31}L]$$

which is 433, so that the expected profit is $Y - 433$. For airline 3 this number is $Y - 352$. When airline 3 is forced to invest in security then the profits for airlines 1 and 2 are each given by $Y - 85$ and profits for airline 3 are $Y - 200$. Hence the profits of each of the three firms are increased when the industry moves for the equilibrium with no investment to a situation with all investing. In fact firms 1 and 2 could profitably pay firm 3 to switch from not investing to investing.

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

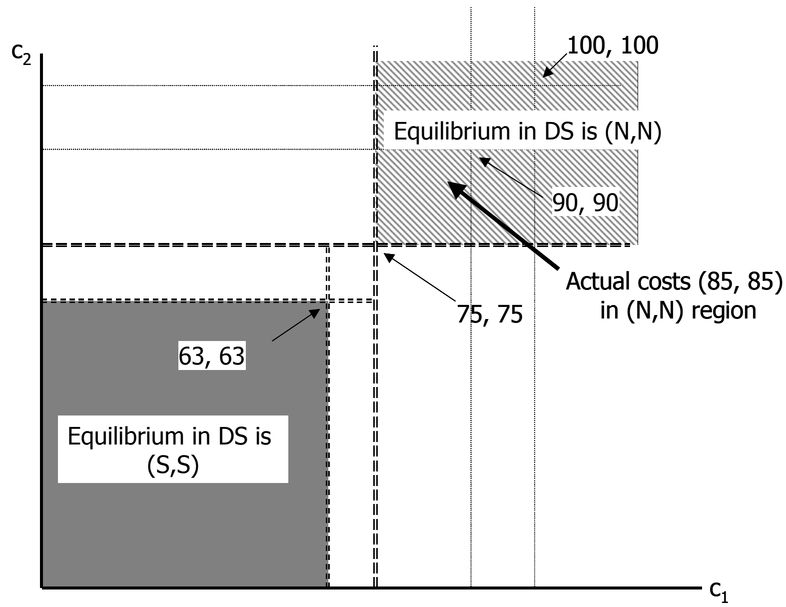

Figure 2: Equilibria for firms 1 and 2 when 3 does not invest.

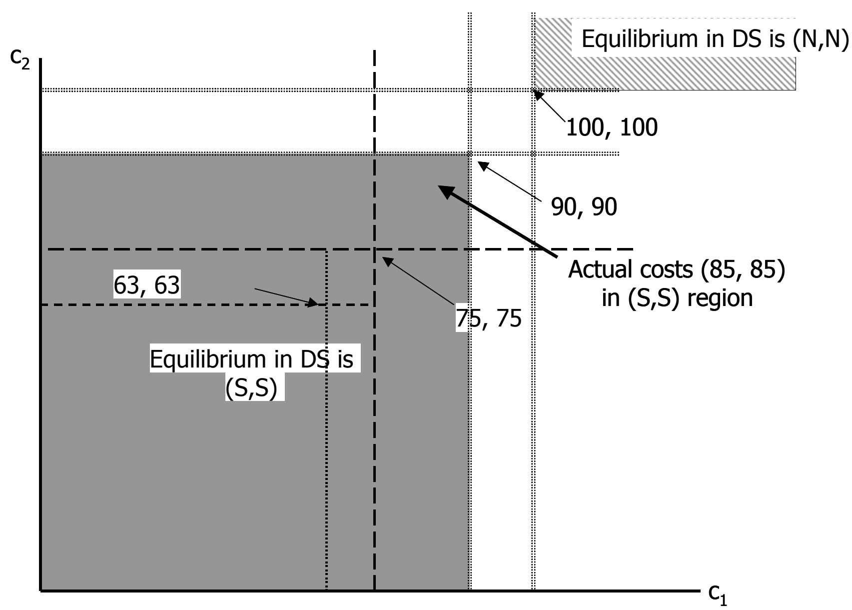

Figure 3: Equilibria for firms 1 and 2 when 3 invests and imposes no externalities. In this case $(85, 85)$ is in the region in which investing is a dominant strategy.

6.2 Cascading

Our model can also give rise to the phenomenon of cascading (see also Dixit [5]), which refers to situation where when one agent changes its policy, this leads another to follow suit. The fact that two agents have changed now persuades a third to follow, and when the third changes policy this creates the preconditions for a fourth to do so, and so on. The analogy with a row of dominoes is compelling: the first knocks down the second, which knocks down the third, and so on. To see how this can happen in our model, suppose that we have a Nash equilibrium at which all airlines choose N and assume in addition we can number firms $1, 2, 3, ...$ so that the following conditions are satisfied:

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

or in general

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

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

In the next two sections we examine how the IDS model applies to a Class 2 problem of full protection (i.e. vaccinations) and a Class 3 problem of positive externalities (i.e. investment in R&D). In both these examples it seems appropriate to assume the probabilities are exogenous since neither problem involves a purposive adversary, such as a terrorist, whose actions are partially determined by the actions taken by different agents in the system.

7 Class 2 Problems: Vaccination

As indicated in Section 2.2, Class 2 problems are ones where an agent who invests in protection obtains complete protection and cannot be contaminated by others. In this section we illustrate the nature of the Nash equilibrium for this class of problems by focusing on deciding whether to be vaccinated or not. We also show the number of individuals who choose to be vaccinated may not be socially optimal.

Catching diseases normally conveys immunity so that you can only catch the disease once: 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.

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, although a terrorist group could play the role played by non-human hosts in the other case. In the absence of deliberate infection by an enemy, we could not normally catch smallpox unless someone else were already infected.

For this example agents may choose to be vaccinated (V) or not to be vaccinated (NV). If you are vaccinated then you will not be infected, so $q_i(K, S) = 0$ whatever the value of K. Define p to 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). Let q denote the chance of catching the disease from a non-human source and infecting another susceptible person. It is only possible to catch the disease once, so that $\alpha = 1$. Y is person i's initial income or welfare, the reference point from which welfare changes are measured.

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 chooses to be vaccinated, the payoffs are the initial wealth Y minus the expected loss from an infection from the background pL, minus also the expected loss 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 $pL + (1-p)\,qL < c$ then (NV, NV) is the equilibrium.

So as the cost of vaccination rises, we have equilibria with both people being vaccinated, one being vaccinated, and neither being vaccinated. 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. This structure persists as we consider situations with more people.

For the case of full protection one can show that there will only be one Nash equilibrium for the reasons indicated in Section 3. The same type of conditions as in the airline security example of the previous section hold for determining when the Nash equilibrium is also socially optimal. If the Nash equilibrium is (V, V) then this strategy maximizes social welfare. For all other cases if the costs of vaccinating are sufficiently low then it will be socially optimal to vaccinate more individuals. In the above 2 person simplified example, consider the case where both (N, V) or (V, N) are equilibria. In this case if the other person were forced to be vaccinated, then the total costs would be $Costs(V, V) = 2c < Costs(N, V) = c + pL + (1-p)\,qL$ by the nature of the cost inequality for(N, V) or (V, N) to be a Nash equilibrium.

8 Class 3 Problems: Investment in R&D

Class 3 problems are ones where a firm that invests in knowledge produces positive externalities for others in the system. We illustrate this type of problem by focusing on the case where two competitive firms in an industry are considering whether or not to invest money in R&D, a topic on which there is an extensive literature (see e.g. Dasgupta and Stiglitz [2] and [3], Dixit [4] and Grossman and Shapiro [7]). Assume that firm i can invest in R&D at a cost of c_i. This generates a payoff of G with probability p_i. There is, in addition, a chance p_j that another firm j invests and succeeds, in which case the information it gains reaches firm i. If I stands for investing and N for not investing then the payoff for the two by two case is

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

1/2	I	N
I	$Y-c_1+p_1G+(1-p_1)\,p_2G,\ Y-c_2+p_2G+(1-p_2)\,p_1G$	$Y-c_1+p_1G, Y+p_1G$
N	$Y+p_2G, Y-c_2+p_2G$	Y, Y

Here if neither invests then there is no chance of either getting the information and so both their payoffs are their initial income Y. If firm one invests and two does not, then the payoff to the investor is $Y-c_1+p_1G$, income net of the cost of investing plus the expected gain from the investment. The payoff to the non-investor here is $Y+p_1G$, income plus the expected gain as the information is transferred to it from the successful investor. Finally if both invest then firm i has a payoff of $Y-c_i+p_iG+(1-p_i)\,p_jG$, 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, I) is a dominant strategy if and only if

$$c_i < p_i\,[1-p_j]\,G \tag{23}$$

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

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 know already from (23) that (I, I) is a Nash equilibrium if

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

Similarly (I, N) is a Nash equilibrium if

$$p_1G > c_1 \text{ and } c_2 > Gp_2\,[1-p_1]$$

and (N, I) is an equilibrium if

$$p_2G > c_2 \text{ and } c_1 > Gp_1\,[1-p_2]$$

We can now look at the plane with c_1 and c_2 as its axes, position the other parameters on this and analyze when (I, I), (N, I), (I, N) and (N, N) are Nash equilibria. In Figure 4 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

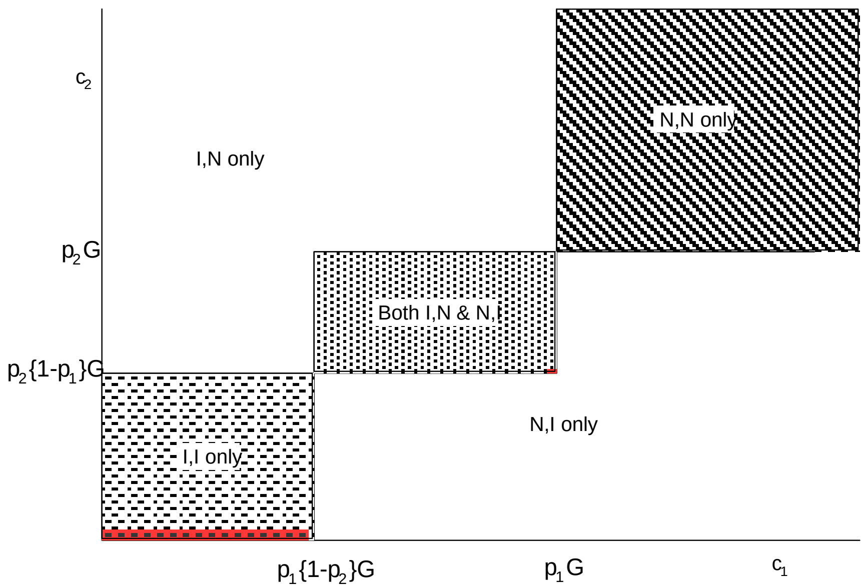

Figure 4: Equilibria of the R&D game as a function of costs c_2 and c_2 and other parameters.

choose to invest and in the upper right region the only equilibria are those where neither chooses to invest. Between these regions is one where there are two possible outcomes, (N, I) and (I, N), and to the upper left the only possible outcomes are (I, N) and to the lower right (N, I). If both firms are identical then Figure 4 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. There is only a single Nash equilibrium so there is no possibility of tipping or cascading as in Class 1 problems.

Figure 4 enables us to examine conditions when the Nash equilibrium maximizes industry expected profits. From the proof of proposition 3 we again know that investing Pareto dominates not investing when

$$c_i < p_i L_i + q_i (\emptyset, N) - \alpha p_i q_i (\emptyset, N)$$

which in the current context reduces to

$$c_i < p_i G + p_j G (1 - p_i)$$

In figure 4 this means that all investing is Pareto efficient whenever costs c_1 and c_2 are less than $p_1 G + p_2 G (1 - p_1)$ and $p_2 G + p_1 G (1 - p_2)$ respectively. This includes the areas where (I, I) and (I, N) and (N, I) are equilibria, and some of the area in which (N, N) is an equilibrium.

9 Conclusions

Interdependence is a widespread phenomenon with risk-management decisions: airlines, electric utilities, public health and R&D amongst others are fields in which the risk that I face depends on what you choose, and vice versa. We have specified three classes of IDS problems and developed a general framework for analyzing them from a game-theoretic perspective. We have identified minimum critical coalitions, the smallest coalitions that can tip an equilibrium, given conditions for them to exist, and shown that they represent the least expensive way of changing from an inefficient to an efficient equilibrium.

An interesting feature of Class 1 problems, where investment in protection only provides partial protection, is the possibility of tipping and cascading. 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 on airline security, a change of strategy from N to S by one airline leads the other two airlines to also invest in prevention. We also used the example to illustrate cascading, where a change of strategy by one agent causes the second to change which induces a third to invest in prevention until all parties have changed their strategy, a classical "domino effect".

The equilibria for IDS problems are often inefficient because of negative external effects. The social return to an investment in protection, in R&D and/or in infection-prevention, is greater than the private return and this can lead to under-investment. The policy implications are interesting: it may be that the private sector through some coordinating mechanism (e.g. a trade association) or the government can identify those "influentials" or "opinion leaders" who form a MCC and persuade them to change their positions. As noted in our illustrative airline security example, the tax needed to influence the minimum critical coalition is much less than that needed to influence all players.[6]

References

[1] Bulow Jeremy I., John D. Geanakoplos and Paul D. Klemperer 1985. "Multimarket Oligopoly: Strategic Substitutes and Complements". *Journal of Political Economy*, vol 93, no 313, 488-511.

[2] Dasgupta Partha and Joseph E. Stiglitz 1981. "Resource Depletion under Technological Uncertainty." *Econometrica* 49, 1: 85-104

[6] In [13] we examine private and/or public sector policy interventions that could be used to correct the underinvestment. These include taxes, subsidies, regulations, third party inspections and the use of associations and other coordinating mechanisms. Lakdawalla and Zanjani [14] also investigate ways in which the public sector can be involved in reducing the negative externalities.

[3] Dasgupta Partha and Joseph E. Stiglitz 1980. "Uncertainty Industrial Structure and the Speed of R&D." *The Bell Journal of Economics* 11, 1: 1-28.

[4] Dixit Avinash K. 1988. "A General Model of R & D Competition and Policy." *The RAND Journal of Economics* 19, 3: 317-326.

[5] Dixit Avinash K 2002. "Clubs with entrapment." Available at www.princeton.edu/~dixitak/home

[6] Gladwell Malcolm 2000. *The Tipping Point* Little Brown and Co.

[7] Grossman Gene M. and Carl Shapiro 1987. "Dynamic R & D Competition." *The Economic Journal* 97, 386: 372-38

[8] Heal Geoffrey 1994. "The Formation of International Environmental Agreements". Pages 301-322 of C. Carraro (ed) *Trade Innovation and the Environment*, Kluwer.

[9] Heal Geoffrey and Kunreuther Howard.(in press). "IDS Models of Airline Security" *Journal of Conflict Resolution.*

[10] Heal Geoffrey. "The Blackout of 03" Financial Times August 26 2003, available at http://www-1.gsb.columbia.edu/faculty/gheal/

[11] Kearns, Michael. Personal Communication, March 2003, Department of Computer and Information Science, University of Pennsylvania. www.cis.upenn.edu/~mkearns/

[12] Keohane Nathaniel O. and Richard J. Zeckhauser 2003. "The Ecology of Terror Defense.". *Journal of Risk and Uncertainty*, Special Issue on Terrorist Risks, Vol 26 No. 2/3 (March/May): 201-229.

[13] Kunreuther Howard and Geoffrey Heal 2002. "Interdependent Security." *Journal of Risk and Uncertainty*, Special Issue on Terrorist Risks, Vol 26 No. 2/3 (March/May): 231-249.

[14] Lakdawalla Darius and George Zanjani (in press) "Insurance, Self Protection and the Economics of Terrorism." *Journal of Public Economics*

[15] Sandler, Todd, 2003. Collective Action and Transnational Terrorism. School of International Relations, University of Southern California.

[16] Schelling, Thomas 1978. *Micromotives and Macrobehavior.* New York: Norton

[17] Watts Duncan J. 1999. *Small Worlds.* Princeton University Press