6-2007

Modeling Interdependent Risks

Geoffrey Heal

Howard Kunreuther
University of Pennsylvania

Follow this and additional works at: https://repository.upenn.edu/oid_papers

Part of the [Aviation Safety and Security Commons](#), [Defense and Security Studies Commons](#), and the [National Security Law Commons](#)

Recommended Citation

Heal, G., & Kunreuther, H. (2007). Modeling Interdependent Risks. *Risk Analysis*, 27 (3), 621-634.
<http://dx.doi.org/10.1111/j.1539-6924.2007.00904.x>

This paper is posted at ScholarlyCommons. https://repository.upenn.edu/oid_papers/54
For more information, please contact repository@pobox.upenn.edu.

Modeling Interdependent Risks

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 suboptimal for everyone. We model this as the Nash equilibrium of a game and give conditions for such a suboptimal 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 suboptimal to the optimal equilibrium. We illustrate these results by reference to airline security and the control of infectious diseases via vaccination.

Keywords

Cascading, critical coalition, interdependence, Nash equilibrium, security, terrorism, tipping

Disciplines

Aviation Safety and Security | Defense and Security Studies | National Security Law

Modeling Interdependent Risks

Geoffrey Heal and Howard Kunreuther*

August 2005 Revised August, November 2006

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 for everyone. 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 and the control of infectious diseases via vaccination.

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

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

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 call this risk class **interdependent security (IDS) problems** and use game-theoretic models to characterize their Nash equilibria. What we are doing here is looking at how individuals and firms manage risks where there is some likelihood that even if they have invested in protective measures they can be harmed as a consequence of others not following suit. We are studying what outcomes result when all the people who are linked in such a system make independent decisions but are aware of the choices made by others. Using game theoretic models we analyze the Nash equilibria and how to improve them if they are not efficient.

The risks of terrorism are typically interdependent: the potential consequences to a firm depend not only on its own choice of security investments but also on the

*Heal: Columbia Business School, gmh1@columbia.edu and Kunreuther: The Wharton School, University of Pennsylvania, kunreuther@wharton.upenn.edu. We are grateful to Vicki Bier and a referee for valuable comments. Support from NSF grant number CMS-0527598 and the Wharton Risk Management and Decision Processes Center is also gratefully acknowledged.

actions of other agents. Failures of a weak link in an interdependent system can have devastating impacts on all parts of the system. These *negative externalities* are an important feature of the IDS problems that will be examined in more detail in this paper.¹ 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, and as may be the case with avian flu.

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 its risks due to 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 they can no longer justify incurring these costs. In equilibrium, all actors may consequently fail to invest in strict security measures.

The explosion of Pan Am 103 over Lockerbie, Scotland, in December 1988 illustrates this point. Terrorists in Malta 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. Unless PanAm inspected all transferred bags (which until recently was true only of El Al) there was no way it could have avoided this crash.

1.1 The Nature of IDS Problems

Before laying out our model in detail, we begin with some comments on how this class of problems relates to the more general literature on games and strategic behavior. At the core of an IDS problem is a stochastic external effect: the possibility of being affected outside of the market by others in the system. That is the essence of the PanAm 103 case mentioned above, in which a failure to take protection by one airline imposed a massive risk on other airlines, even if they invested in security precautions. External effects are of course not new, though the literature on stochastic externalities is limited. What is interesting is that a class of apparently rather different problems involving stochastic externalities all have the same, rather specific, mathematical structure, so that a common framework emerges which is what

¹The externalities associated with IDS can also be positive as we indicate for the problem of optimal investment in Research and Development (R&D).

we call the IDS framework. This encompasses issues as apparently diverse as airline security, the management of infectious diseases, security in computer networks, and the payoffs from investing in R&D. All of these cases have a similar structure. It is the differences in the details that lead to different outcomes. The principal purpose of this paper is to demonstrate similarities and difference between these problems. The results presented here also extend and generalize findings from earlier papers [16], [8].

At the core of our analysis is a result that we spell out in detail in another more technical paper [10], which states that there is a close connection between two literatures that have previously been seen as quite unconnected, papers on tipping and those on supermodularity and strategic complementarity. What the results here illustrate is that supermodularity of a game is close to being a sufficient condition for the equilibria of the game to exhibit the tipping property discussed by Schelling [21], [22] and many others since his seminal papers. In non-technical terms, the essence of cases such as the PanAm 103 crash is that the failure of one agent to attend to security issues reduces the effectiveness of the measures that other agents may take to improve their security, and so reduces their incentive to take such measures. The return to one agent's investment in security depends on the actions of other agents. If others do invest, then investing is more attractive, and if they do not, then it is less. This means that the underlying game exhibits strategic complementarity or supermodularity (Bulow et al. [2], Milgrom and Roberts [19]). We show below that there are two natural categories of IDS phenomena: airline security issues show strategic complementarity, whereas fully effective vaccines and R&D investments exhibit strategic substitutability.

1.2 Outline of the Paper

The next section sets out each of the three classes of IDS problems by indicating the nature of their Nash equilibria and the implications of this behavior for private and social welfare. In section 3 we present a general model, of which all these classes are special cases, and derive results from this general model that can be applied to each of the classes of IDS problem. These results specify conditions for tipping and characterize the smallest coalition needed to tip an equilibrium from an inefficient to an efficient outcome. In our initial treatments the probabilities of a direct loss to an agent from an external source (e.g. a terrorist) is treated as exogenous. In section 4 we extend the analysis to the case of endogenous probabilities of direct losses. For example, the likelihood that a terrorist will attack you directly will now depend on whether or not other agents have taken action to protect themselves. Finally in sections 5 and 6 we spell out in detail the applications of the IDS framework re two important types of problems that relate to terrorism: airline security and vaccination against infectious diseases.

2 Classes of IDS Problems

Interdependent security (IDS) problems have one common characteristic—the decision by one agent as to whether or not to incur an investment cost impacts the welfare of other agents, and thus affects their economic incentives to take action themselves. This section categorizes three classes of IDS problems based on differences in their Nash equilibria.

2.1 Class 1: Partial Protection with Negative Externalities

Here an agent’s decision to invest in protection so as to reduce its own risks also decreases the risks experienced by others. In other words, the more that agents 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 containing a bomb. The more airlines invest in baggage security, the greater the reduction in the risk experienced by everyone 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 invest in protective measures herself. The decision by electric utilities to take steps to reduce the likelihood of a power failure is 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 [12]).

2.1.1 Nash Equilibria

As we show more formally in the next section, this class of problems can have one or multiple Nash equilibria. For the case where there are two equilibria: either all agents invest in security or none of the agents want to do this, there is then the possibility of tipping or cascading. In other words, inducing some agents to invest in protection 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 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. Society will be better off as well.

When there is only a single Nash equilibrium, the investment choices by each of the individual agents will also be socially optimal in some situations. The most obvious case is where the costs of protection are sufficiently low so that each agent wants to invest in protection even when all the other agents did not incur these costs. If investment costs 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 sufficiently high 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 3–4, and an illustrative example with respect to airline security is presented in Section 5.

2.2 Class 2: Complete Protection with Negative Externalities

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 vaccination 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 (Kunreuther and Heal [17]).

2.2.1 Nash Equilibrium

For Class 2 problems there is only one type of 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. Consider agent i . As more other agents invest in protection there is a reduction in the negative externalities that impact on i . This translates into a lower probability of agent i suffering a loss which means that it now has a lower expected benefit of investing in protection

2.2.2 Private and Social Welfare

As in Class 1 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 discuss the vaccination problem with complete protection in Section 6 of the paper.

2.3 Class 3: Investments with Positive Externalities

For this class of problems an investment by one agent creates positive externalities, which make 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 higher the probability 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 the 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 cost 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 do not discuss the R&D problem in more detail in this paper - for details see Heal and Kunreuther [9].

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, and specify conditions for multiple equilibria. We focus on the case where one equilibrium involves investment in security by all agents while the other involves no

investment by any agent. 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 and show that they are generically unique and identical to the (unique) cheapest critical coalition. Strategic complementarity and substitutability (Bulow Geanakoplos and Klemperer [2]) and supermodularity (Milgrom and Roberts [19]) lie at the heart of some of the phenomena that we study.

We consider a total of A interdependent risk-neutral agents indexed by i . Each agent is characterized by parameters p_i , L_i , and c_i . Here p_i is the probability that agent i 's actions lead to a direct loss L_i . By investing in loss-prevention at a cost of c_i , agent i avoids a direct loss with certainty. Each agent i has a discrete strategy, X_i , that can take as values either S or N representing investing and not investing respectively. If agent i incurs a direct loss, then this may also affect other agents' outcomes. If agent i does not incur a direct loss then it will have no negative impact on others. We call the loss (or in some cases gain) to others an indirect impact. More specifically $q_i(\{K\}, X_i)$ is the expected indirect loss to agent i when it follows strategy X_i and the agents in the set $\{K\}$ are the only ones investing in loss-prevention. 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 in particular if everyone other than i invests in prevention, then i cannot suffer indirect impacts. That is if $\{K\} = \{1, 2, \dots, i-1, i+1, \dots, A\}$ 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 loss 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 $p_i L + q_i(\{K\}, N)$, 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

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

²For discussions of tipping and cascading in the literature, see Schelling [21] [22], Dixit [4], Watts [23] and Gladwell [6]. Easterly [5] provides a brief history of the literature on tipping and tests the prediction of the model for racial segregation.

or

$$c_i^* (\{K\}) \triangleq p_i L_i + (1 - \alpha p_i) q_i (\{K\}, N) - q_i (\{K\}, S) \quad (2)$$

where $c_i^* (\{K\})$ in equation (2) 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 [8] [16] where the indirect risk is the same whether agent i invests or not: $q_i (\{K\}, N) = q_i (\{K\}, S)$ and $\alpha = 1$ so that

$$c_i^* (\{K\}) \triangleq p_i (L_i - q_i (\{K\}, N)) \quad (3)$$

Since the probability of an indirect loss, q_i decreases as another agent is added to $\{K\}$, it follows that $c_i^* (\{K\})$ increases in $\{K\}$. In words as more agents invest, then the expected indirect loss falls and the maximum cost at which investment is justified rises, with $c_i^* (\emptyset) < c_i^* (\{A/i\})$,³ the latter being the critical cost when all agents other than i are investing. In this case the game is supermodular (see Milgrom and Roberts [19]). 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 $\alpha = 1$. This implies that if agent i has invested in protection there is no chance of i suffering an indirect loss from others so that

$$c_i^* (\{K\}) \triangleq p_i L_i + [1 - p_i] q_i (\{K\}, N) \quad (4)$$

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 (right hand side) of the equation defining $c_i^* (\{K\})$ is negative for Class 1 problems and positive for Class 2 problems. In Class 3 problems, such as determining whether to invest in R&D, those who undertake investments may make a discovery and there is the possibility of a positive benefit from investment B_i . Hence q_i is also positive. In addition, we assume that the likelihood of benefiting from someone else's discovery is independent of whether agent i invests in R&D themselves so that $q_i (\{K\}, N) = q_i (\{K\}, S)$.

Those who invest in R&D and make a discovery provide spillover benefits to others who may not have invested in R&D but who can use the discovery, so that

$$c_i^* (\{K\}) = p_i [q_i (\{K\}, S) - B_i] \quad (5)$$

As in Class 2 problems $c_i (\{K\})$ **decreases** with $\{K\}$, but for different reasons—agents benefit from others' investments.

We now investigate properties of the Nash equilibria of this system. We consider only equilibria in pure strategies.

Definition 1 *A Nash equilibrium for the above class of problems is a set of pure strategies $X_1, \dots, X_A$ such that (1) $X_i = S$ for all $i \in \{K\}$ (which may be empty), (2) if $X_i = S$ then $c_i^* (\{K\}) > c_i$ and (3) if $X_i = N$ then $c_i^* (\{K\}) < c_i$ and (4) if $c_i^* (\{K\}) = c_i$ then i is indifferent between S and N .*

³ $\{A/i\}$ means the set A without the element i .

Theorem 1 *For problems in Class 1, a Nash equilibrium exists.*

Proof. An adaptation of the arguments in Kearns [14]: this theorem also follows from results in Milgrom and Roberts [19]. ■

There may be equilibria where all agents invest in loss-prevention, those where none do, and asymmetric pure strategy 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:

Theorem 2 *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 $(N, N, \dots, N)$ and $(S, S, \dots, S)$ are Nash equilibria, then $(S, S, \dots, S)$ Pareto dominates $(N, N, \dots, N)$.*

Proof. See appendix. ■

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

3.1 Tipping

In this section we explore the possibility that a game with two (or more) equilibria may be tipped from an inefficient to an efficient equilibrium by a change in the strategy choices of a small number of players. We show that this is possible for class 1 IDS problems, those that exhibit strategic complementarity. In this case a decision by, for example, one airline to invest improves the incentives for others to follow suit. It should be intuitively clear that if the mutual reinforcement effect is strong enough, then this could begin a process that will lead all agents to change.

Definition 2 *Let $X_i = N \forall i$ be a Nash equilibrium. A critical coalition CC for this equilibrium is a set $\{M\}$ of agents such that if $X_i = S \forall i \in \{M\}$ then $c_j^*(\{M\}) \geq c_j \forall j \notin \{M\}$.*

In words, a critical coalition is one with the property that if all of its members invest then for all other agents the best response is also to invest.

Definition 3 *A minimum critical coalition MCC is a critical coalition of which no subset is also a critical coalition.*⁴

We are interested in critical coalitions in which there is no redundancy and all agents in the coalition are needed to tip the other agents.

⁴See Heal [13] for an earlier use of the idea of a minimum critical coalition in a different context, unconnected with tipping.

Definition 4 *A smallest critical coalition SCC is a minimum critical coalition with the property that no other critical coalition contains fewer members.*

There can be several minimum critical coalitions - we are interested in the smallest one. Define

$$q_i^j(\{K\}, N) = q_i(\{K/j\}, N) - q_i(\{K\}, N) \geq 0, \quad j \in \{K\} \quad (6)$$

This represents 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 A1: $q_i^j(\{K\}, N)$ is independent of i : $q_i^j(\{K\}, N) \triangleq q^j(\{K\}, N) \forall i$

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

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

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:

Assumption A3 : The ranking of agents $j = 1..A$ by the indirect effects they generate $q^j(\{K\})$, is independent of $\{K\}$

This says in intuitive terms that if agent 1 creates the largest negative externalities when agents in the set $\{K\}$ are investing in loss-prevention, then agent 1 creates more externalities than any other agent no matter who is in the set investing in loss prevention.

Theorem 3 *Let $X_i = N \forall i$ be a Nash equilibrium. If a smallest critical coalition exists for this equilibrium then for some integer k it consists of the first k agents when agents are ranked in decreasing order of $q^j(\emptyset)$.*

Proof. See the Appendix ■

Corollary 1 *There is a smallest critical coalition only if $\alpha > 0$, i.e., if there is some degree of non-additivity of damages.*

Corollary 2 *A smallest 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 *SCC* could not be unique is if $q^j(\{K\}) = q^f(\{K\})$ for some $j, f < L$, which is ruled out by the assumption. ■

These results imply that a *SCC* is easily characterized and that in general there is only one *SCC*, if the terms $q^j(\{K\})$ will differ. Note that assumption A2 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 $\alpha = 0$.

The key policy implications of our results on critical coalitions is that an equilibrium where no one invests in security may be converted to one with full investment by persuading a subset of the agents to change their policies. Tipping assures us that it is only necessary to persuade the subset to alter their behavior to convince all agents to invest in protection. The least expensive way of changing equilibrium is therefore likely to involve providing incentives for a critical coalition to change its behavior, which will then tip the entire system.⁵

4 Endogenous Risks of Direct Losses

In the above analyses the direct 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 [16]. 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 [20]. Keohane and Zeckhauser [15] also consider the implication of endogenous terrorist risks, focusing on ways to controlling the stock of terrorist capital and curbing the flow into the terrorist organizations. Bier, Oliveros and Samuelson [1] develop a strategic model for determining how much a defender should allocate to each location to defend against possible attacks, knowing that the likelihood that a terrorist will attack a specific location decreases as more resources are allocated to protecting it.

To examine the impact of endogenous risks on Nash equilibria we focus here on Class 1 problems and the case of airline security. We assume that the risk faced by an airline that does not invest in stricter inspections increases as the fraction of other 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 η 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, $p_i(\eta)$, are now increasing in η .

Now return to equation (3), 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 $q_i(\{K\}, N) = q_i(\{K\}, S)$ and $\alpha = 1$ so

⁵Heal and Kunreuther [10] provide conditions for finding the smallest set of agents that can tip all the other agents.

that

$$c_i^* (\{K\}) \triangleq p_i (\eta) [L_i - q_i (\{K\}, N)] \quad (7)$$

In this expression $q_i (\{K\}, N)$ depends on η since the likelihood of airline i being impacted by others depends how many airlines are investing in security. To understand how a change in η will affect $c_i (\{K\})$, take the first difference of the right hand side of (7) with respect to η :

$$\frac{\Delta c_i (\{K\})}{\Delta \eta} = [L_i - q_i (\{K\}, n)] \frac{\Delta p_i}{\Delta \eta} - p_i \frac{\Delta q_i}{\Delta \eta} \quad (8)$$

Here Δ is the difference operator and $\frac{\Delta p_i}{\Delta \eta} > 0$ by assumption, and the coefficient

associated with this term is the difference between the direct and expected indirect losses, which we assume to be positive. The first term on the RHS of (8) is therefore positive under this assumption. The term $\frac{\Delta q_i}{\Delta \eta}$ 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 3 on tipping is thus relevant to a model with endogenous probabilities. With the above assumptions, $\frac{\Delta c_i^* (\{K\})}{\Delta \eta} > 0$ in (8) 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.

5 Class 1 Example: Airline Security⁶

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 the related phenomenon of cascading through numerical examples coupled with a geometric framework to provide intuition for these results. There are 2 separate airlines. Let r_{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

⁶This section is based on material that appeared in [8].

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 costs per trip are just their investment costs, c_i . If A_1 invests and A_2 does not, then A_1 incurs an investment cost c_1 plus the expected loss from a bomb transferred from A_2 that explodes on A_1 (i.e., $r_{21}L$), while A_2 has an expected loss from a bomb loaded and exploding on its plane, $r_{22}L$. If neither invests then A_1 has an expected loss from a bomb loaded and exploding on its own plane $r_{11}L$ minus the expected loss from a bomb transferred from A_2 that explodes on A_1 (i.e., $r_{21}L$) conditioned on there being no explosion from a bomb loaded by A_1 itself $(1 - r_{11})$. A_2 's payoff is determined in a similar fashion.

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

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 theorem 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

$$c_1 < r_{11}L \text{ and } c_1 < r_{11}[1 - r_{21}]L \quad (9)$$

The condition that $c_1 < r_{11}L$ is clearly what we would expect from a single airline operating on its own. The tighter condition that $c_1 < r_{11}[1 - r_{21}]L$ 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 $c_i < r_{ii}L$ and is a dominant strategy if $c_i < r_{ii}L(1 - r_{ji})$ where i and j are 1 or 2. (N, N) is a Nash equilibrium if $c_i > r_{ii}L(1 - r_{ji})$ and a dominant strategy if $c_i > r_{ii}L$. From these inequalities we note that (S, S) and (N, N) are both Nash equilibria if $r_{ii}L(1 - p_{ji}) < c_i < r_{ii}L$: this is consistent with theorem 2, as $r_{ii}L(1 - r_{ji}) = c_i(\emptyset)$ and $r_{ii}L = c_i(A - i)$. Finally if $c_1 > r_{11}L$ but $c_2 < r_{22}L(1 - r_{12})$ 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 [16].

Figure 1 enables one to determine when the Nash equilibrium also is socially optimal. From the proof of Theorem 2 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. To be precise we know from the proof of Theorem 2 that (S, S) Pareto dominates (N, N) whenever

$$c_i < p_i L_i + (1 - \alpha p_i) q_i(\emptyset, N) \quad (10)$$

In the present context this simplifies to

$$c_i < r_{ii} L + L(1 - r_{ii}) r_{ji} \quad (11)$$

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 $[0, r_{11}L]$ and $[0, r_{22}L]$. 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.

5.1 Tipping

Consider three airlines, and let $r_{11} = r_{12} = r_{13} = r_{21} = r_{22} = r_{23} = 0.1$, $r_{31} = r_{32} = 0.3$, $r_{33} = 0.2$, $L = 1000$ 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

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

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

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

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 $c_i > 63, i = 1, 2$.

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 (9), which gives a critical cost level of 90, meaning that investment will now be a dominant strategy when the cost is less than 90. Since $c_1 = c_2 = 85$, 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.⁷ When firm 3 does not invest, as in figure 2, not investing is a dominant strategy for the other two 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, as reflected 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 $(85, 85)$ representing the investment costs of 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

$$r_{11}L + (1 - r_{11})[r_{21}L + (1 - r_{21})r_{31}L]$$

which is 433, so that its expected profit is $Y - 433$. For airline 3 this value 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 from 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.

5.2 Cascading

Our model can also give rise to the phenomenon of cascading (see also Dixit [4]), which refers to a situation where when one agent changes its policy, this leads another to

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

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

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.

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 $1, 2, 3, \dots$ so that the following conditions are satisfied:

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

or in general

- When $1, 2, 3, \dots, J$ 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.

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 $c_1 = 95$ as before but $c_2 = 85$. 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 $c_1 < r_{11}L = 100$, 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.

6 Class 2 Problems: Vaccination⁸

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 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

⁸For more details on the vaccination problem see [11].

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. Because of this, the incentive that any agent faces for vaccination is reduced as more other people are vaccinated: their vaccination reduces the risk to which he or she is exposed and so reduces the benefits of vaccination. Hence we have a case of strategic substitutability.

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,⁹ so $q_i(K, S) = 0$ whatever the value of K . Define r 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 $\alpha = 1$. 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	$-c, -c$	$-c, -rL$
NV	$-rL, -c$	$-rL - (1 - r)qL, -rL - (1 - r)qL$

If both are vaccinated then each incurs a the cost of vaccination, $-c$. If only one is vaccinated then her cost is $-c$, and the other has an expected loss $-rL$: the latter person runs no risk of infection from the vaccinated person who by assumption cannot transmit the disease.

In the case in which neither chooses to be vaccinated, an individual has an expected direct loss of $-rL$, plus an expected indirect loss from infection by the other person $-qL$, which only occurs if you have not already been infected, an event with probability $(1 - r)$. From this payoff matrix it is clear that:

1. When $c < rL$, (V, V) is a Nash equilibrium.
2. For $rL < c < rL + (1 - r)qL$, both (N, V) or (V, N) are equilibria, and

⁹In fact many vaccines are not fully effective. For example, most vaccines against influenza convey only partial immunity to the disease.

3. For $rL + (1 - r)qL < c$ 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 (rL), and the expected loss from infection if she is not $(r + (1 - r)q)L$. Here $(r + (1 - r)q)$ is the probability of infection if neither is vaccinated. This same structure persists as we consider situations with more people.

If the Nash equilibrium is (V, V) then this strategy maximizes social welfare because $c < rL$. 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 which case $rL < c < rL + (1 - r)qL$. In this case if the other person were forced to be vaccinated, then the total costs would be $Costs(V, V) = 2c < Costs(N, V) = c + rL + (1 - r)qL$ by the nature of the cost inequality for (N, V) or (V, N) to be a Nash equilibrium. So the outcome would not be socially optimal.

7 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. This can lead to strategic complementarity or substitutability. We have specified three classes of IDS problems and developed a general framework for analyzing them from a game-theoretic perspective, drawing on the complementarity or substitutability between agents' strategies. We have illustrated the challenges in developing protective strategies for dealing with terrorism that illustrate Class 1 problems (airline security) and Class 2 problems (vaccinations).

An interesting feature of Class 1 problems, where investment in prevention only provides partial protection and where investments are strategically complementary, 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 externalities. The social return to an investment in protection and in disease-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) can induce all firms in the industry to

invest in cost-effective preventative measures. Or the government can identify those “influentials” or “opinion leaders” who form a minimum critical coalition and persuade them to change their positions, leading others to then adopt similar measures. 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.¹⁰

8 Appendix

Proof of Theorem 2

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 $(S, S, \dots, S)$ equilibrium is equivalent to

$$c_i < p_i L_i + (1 - \alpha p_i) q_i(\emptyset, N) \quad (12)$$

where the left hand side (LHS) of (12) reflects the costs to each agent i if all agents invest in prevention and the RHS of (12) is the cost to agent i if no-one invests in prevention, since $q_i(\{K\}, S) = 0$ when $\{K\} = \{A\}$. The existence of both $(S, S, \dots, S)$ and $(N, N, \dots, N)$ as equilibria implies that

$$c_i < c_i^*(\{A/i\}) = p_i L_i + (1 - \alpha p_i) q_i(\{A/i\}, N) - q_i(\{A/i\}, S) \quad (13)$$

The RHS of (13) is less than the RHS of (12) because $q_i(\emptyset, N) > q_i(\{A/i\}, N)$ and $q_i(\{A/i\}, S) \geq 0$ so that (13) implies (12). This completes the proof of the proposition. ■

Proof of Theorem 3

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

$$\Delta c_i^j = \{q_i(j, N) - q_i(j, S)\} + p_i(L_i - \alpha q_i(j, N)) - \quad (14)$$

$$\{q_i(\emptyset, N) - q_i(\emptyset, S)\} - p_i(L_i - \alpha q_i(\emptyset, N)) \quad (15)$$

$$= (1 - \alpha) \{q_i(j, N) - q_i(\emptyset, N)\} + \{q_i(\emptyset, S) - q_i(j, S)\} \quad (16)$$

Using A3, we see that for $\{K\} = \{1, 2, 3, \dots, k\}$ 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

$$\sum_{j=1}^{j=k} \Delta c_i^j \geq c_i - c_i(\emptyset) \forall i > k \quad (17)$$

¹⁰In [16] 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 [18] also investigate ways in which the public sector can be involved in reducing the negative externalities.

which can be written as

$$\sum_{j=1}^{j=k} (q^j(\emptyset, S) - (1 - \alpha) q^j(\emptyset, N)) \geq c_i - c_i(\emptyset) \forall i > k \quad (18)$$

By (??) this can be simplified to

$$\alpha \sum_{j=1}^{j=k} q^j(\emptyset) \geq c_i - c_i(\emptyset) \forall i > k \quad (19)$$

As there can be no smaller coalition that is critical, this completes the proof. ■

References

- [1] Bier Vicki, Santiago Oliveros and Larry Samuelson. (2007) "Choosing What to Protect" . *Risk Analysis* (this Issue)
- [2] Bulow Jeremy I., John D. Geanakoplos and Paul D. Klemperer. (1985) "Multi-market Oligopoly: Strategic Substitutes and Complements". *Journal of Political Economy*, vol 93, no 313, 488-511.
- [3] Dixit Avinash K. (1988) "A General Model of R & D Competition and Policy." *The RAND Journal of Economics* 19, 3: 317-326.
- [4] Dixit Avinash K. (2002) "Clubs with entrapment." Available at www.princeton.edu/~dixitak/home
- [5] Easterly, William. (2004) "Empirics of Strategic Interdependency" February (mimeo).
- [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 and Howard Kunreuther. (2005) "IDS Models of Airline Security" *Journal of Conflict Resolution* 41:201-17.
- [9] Heal Geoffrey and Howard Kunreuther. (2005) "IDS Models for R&D Problems" Center for Risk Management and Decision Process Working Paper No. 05-11-HK, August.
- [10] Heal Geoffrey and Howard Kunreuther. (2006) "Supermodularity and Tipping" NBER Working Paper #12281, June.
- [11] Heal Geoffrey and Howard Kunreuther. (2005) "The Vaccination Game" Center for Risk Management and Decision Process Working Paper No. 05-10-HK, February.

- [12] Heal Geoffrey. (2003) "The Blackout of 03" Financial Times August 26 2003, available at <http://www-1.gsb.columbia.edu/faculty/gheal/>
- [13] Heal Geoffrey. (1994) "The Formation of International Environmental Agreements". Pages 301-322 of C. Carraro (ed) *Trade Innovation and the Environment*, Kluwer.
- [14] Kearns Michael. Personal Communication, March 2003, Department of Computer and Information Science, University of Pennsylvania. www.cis.upenn.edu/~mkearns/
- [15] 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.
- [16] 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.
- [17] Kunreuther Howard and Geoffrey Heal. (2005) "Interdependencies within an Organization" in B. Hutter and M. Powers (ed) *Organizational Encounters with Risk* (Cambridge University Press).
- [18] Lakdawalla Darius and George Zanjani. (2005) "Insurance, Self Protection and the Economics of Terrorism." *Journal of Public Economics* 89: 1891-1905.
- [19] Milgrom Paul and John Roberts. (1990) "Rationalizability, Learning and Equilibrium in Games with Strategic Complementarities" *Econometrica* 58: 1255-77.
- [20] Sandler, Todd. (2003) *Collective Action and Transnational Terrorism*. School of International Relations, University of Southern California.
- [21] Schelling Thomas. (1971) "Dynamic Models of Segregation" *Journal of Mathematical Sociology* 1:143-86.
- [22] Schelling Thomas. (1978) *Micromotives and Macrobehavior*. New York: Norton.
- [23] Watts Duncan J. (1999) *Small Worlds*. Princeton University Press.