

NBER WORKING PAPER SERIES

SUPERMODULARITY AND TIPPING

Geoffrey Heal
Howard Kunreuther

Working Paper 12281
<http://www.nber.org/papers/w12281>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
May 2006

Heal is at Columbia Business School and Columbia School of International Affairs, gmh1@columbia.edu. Kunreuther is at the Wharton School of the University of Pennsylvania, kunreuther@wharton.upenn.edu. We are especially grateful to Larry Samuelson for assistance with this paper, and also grateful to Doug Bernheim, Vince Crawford and two referees for constructive comments. The views expressed herein are those of the author(s) and do not necessarily reflect the views of the National Bureau of Economic Research.

©2006 by Geoffrey Heal and Howard Kunreuther. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Supermodularity and Tipping
Geoffrey Heal and Howard Kunreuther
NBER Working Paper No. 12281
May 2006
JEL No. C72, D80, H23

ABSTRACT

We model tipping as a game-theoretic phenomenon and investigate the connection between supermodular games, tipping of equilibria and cascading, and apply the results to issues that arise in the context of homeland security and computer security. We show that tipping and cascading can occur in supermodular games and that "increasing differences" is a sufficient condition for tipping. Supermodularity and tipping of equilibria are closely related. We relate our results to Schelling's early work on tipping.

Howard Kunreuther
Operations and Information Management
The Wharton School
University of Pennsylvania
Philadelphia, PA 19104-6366
and NBER
kunreuther@wharton.upenn.edu

1 Introduction

This paper brings together two very different literatures, one on tipping and cascading, and one on supermodularity and strategic complementarity. The idea of tipping was made widely known in economics by Thomas Schelling's pathbreaking work in 1978 [13], although it goes back further as his book indicates.¹ The idea is that a small change in the state of a system can cause a large jump in its equilibrium.² In one of the best-known of Schelling's examples the state variable was the racial composition of a neighborhood, and at the tipping point a small increase in the proportion of non-whites led to a new equilibrium with a sharply lower fraction of whites. Malcolm Gladwell's recent best-seller [5] popularized the idea of tipping and intimated its applicability to a vast range of phenomena. Cascading is a refined form of tipping (see Avinash Dixit [3]) in which the movement from one equilibrium to another occurs through a classic domino effect. A change of strategy by agent one leads agent two to follow suit, and the changes by one and two together lead agent three to follow and so on. There is a step-by-step movement to a new equilibrium, initiated by a change on the part of one agent.

It is remarkable that though tipping and cascading are so widely referenced in the social sciences, they have not previously been modelled in a game-theoretic framework or linked to a broader context (with the exception of Dixit's paper [3]). We argue

¹See also Schelling 1971 [12]. William Easterly 2004 [4] gives a more detailed history of this concept.

²Intuitively it seems that there should be a connection with catastrophe theory, although this point does not seem to have been explored in the literature.

in this paper that tipping and cascading are natural outcomes in a wide variety of problems that exhibit supermodularity.

The literature on supermodularity also goes back to the 1970s, dating to the work of Donald Topkis [14] [15], although it and the literature on tipping have evolved quite separately. The idea of supermodularity was developed further by Xavier Vives [16] and Paul Milgrom and John Roberts [11]: since then it has been used widely in the literatures on game theory and comparative statics. Jeremy Bulow et al. [1] introduced the related idea of strategic complementarity. Supermodularity allows us to identify a class of games for which rather general comparative static results are available, and builds on the idea of increasing differences, which means that the return to a move by one agent can be increased by actions by other agents.

Our work has evolved from research on strategic interdependence in the context of national security. Originally motivated by a desire to understand the impact of interdependence in airline security after 9/11/01, it has evolved to a more general model of how interdependence affects the incentive to invest in protective measures for any kind of network, including electronic networks such as computer networks. (For the national security applications see our papers [6] [9], and for computer network applications see Michael Kearns [8].) One of our early findings was that many networks exhibit a tipping phenomenon with respect to investments in security: for certain values of state variables few agents invest and security is low. A small change can lead to everyone investing with a massive increase in security, a finding that might

have important policy implications. In trying to understand why tipping occurs we inevitably are drawn to supermodularity, as this is a template of a particular type of interdependence that is well-understood.

Our central result is that supermodularity and tipping/cascading are related. In particular, the “increasing differences” property that is linked to supermodularity is a key to tipping and cascading, and indeed is a sufficient condition for tipping. A necessary and sufficient condition for the equilibria of a symmetric game to exhibit the tipping phenomenon is that the game show “sufficient increasing differences.” In the case of symmetric games we give a highly intuitive definition of what it is to show sufficient increasing differences and then use this to describe a simple and intuitive algorithm for constructing the smallest possible tipping set, i.e. the smallest set such that when all its members change strategy then all agents outside the set follow suit. We show similar results for cascading. A key part of the proof is providing a definition of tipping in terms of the properties of Nash equilibria, rather than as the equilibria of dynamic processes as in the generally-used examples.

Intuitively it is clear that there is a connection between strategic complementarity and tipping, and hence between supermodularity and tipping. The essence of strategic complementarity is that one agent making a particular move makes that same move more attractive to other agents. It then follows that this may lead to tipping. However there does not seem to be a formal analysis of this relationship, which is what we are offering here. After developing the theoretical relationship between supermodularity

and tipping, we illustrate our results with the airline security problem that motivated our original research. More specifically we show how the main theorem can be used to identify airlines that play a critical role in the adoption of higher standards of security in a decentralized system. (A numerical application with real airline data can be found in [6].)

2 Tipping and supermodularity

Consider a game with N players $i \in \{1, 2, \dots, N\}$, each choosing a strategy s_i in the discrete strategy space $\{0, 1\}$ and having a utility function $u_i : \{0, 1\}^N \rightarrow R$. We have a natural order on the hypercube $\{0, 1\}^N$ given by the standard vector ordering on R^N , and with this ordering denoted \geq it forms a complete lattice.³

We assume that each agent's payoff function u_i shows increasing differences in the choices of strategies by other agents. Formally this means that if 0_i or 1_i denotes a 0 or 1 in the i^{th} position of the vector S of all strategy choices and S_{-i} denotes the vector of choices of all agents other than i , and for $S'_{-i} \geq S_{-i}$, then for any i

$$u_i(1_i, S'_{-i}) - u_i(0_i, S'_{-i}) \geq u_i(1_i, S_{-i}) - u_i(0_i, S_{-i}) \text{ with strict inequality if } S'_{-i} > S_{-i} \quad (1)$$

This implies that the payoff to agent i from changing from 0 to 1 increases if another agent changes from 0 to 1. This property is implied by supermodularity of

³We use $A > B$ to show that A exceeds B in at least one component and is no less in all and $A \geq B$ to show that it is at least as great in all components.

the functions u_i in the sense

$$u_i(S') + u_i(S'') \leq u_i(S' \vee S'') + u_i(S' \wedge S'')$$

where $S' \vee S''$ and $S' \wedge S''$ denote respectively the least upper bound and greatest lower bounds of the vectors S' and S'' respectively (for details see Milgrom and Roberts [11]).

We shall assume that the game has (at least) two Nash equilibria, $\{0, 0, \dots, 0\}$ and $\{1, 1, \dots, 1\}$. We also assume that these equilibria are Pareto ranked with the latter dominating, as this is the case in the applications that motivate this paper. We shall refer to these on occasions as the efficient and inefficient equilibria. However it is important that none of the propositions to be established below depend in any way on the equilibria being Pareto ranked: they depend only in the increasing differences assumption (1).⁴ A policy-maker will naturally be interested in ruling out the inefficient equilibrium and ensuring an efficient outcome, as in a coordination problem (Vincent Crawford [2], Walter Heller et al. [7]). We study conditions under which it is possible to do this by changing the strategies of a subset of the players. This is the tipping problem: a “tipping set” of agents can by changing their strategies shift the equilibrium from one extreme to the other.

⁴Milgrom and Roberts in Theorem 5 and corollaries establish that a supermodular game has largest and smallest serially undominated strategy profiles and that these are Nash equilibria. The zero and one profiles are undominated, but we have not assumed supermodularity, just increasing differences.

Let T be an arbitrary subset of players. We are going to investigate whether agents in the set T can “tip” the system, i.e. can by changing strategy shift the equilibrium from $\{0, 0, \dots, 0\}$ to $\{1, 1, \dots, 1\}$. To do this we define the T – game as the game above with the additional constraint that for all players in T the only permissible strategy choice is $s_i = 1$. If the T – game has as its only equilibrium $\{1, 1, \dots, 1\}$ then we say that T is a tipping set or T – set. The key point here is that when agents in T choose strategy 1, this is also the best response for all agents not in T .

A set is a *minimal* T – set if it is a T – set and no subset is. It is a *smallest* T – set if it is a T – set and no other T – set has fewer members. Clearly if T is a T – set then getting the members of T to adopt strategy 1 will rule out the inefficient equilibrium: members of the set T can tip the equilibrium, can force the system to the efficient outcome. If T is a small subset of N then this can be an important policy tool.

Below we characterize these various types of T – sets and show that in certain cases the smallest T – set can be formed by a simple algorithm, in which we rank agents by a very natural characteristic and then pick the first $K \leq N$ in this ranking. Intuitively the characteristic is a measure of the changes in others agents’ payoffs that result when an agent changes her strategy from 0 to 1. It is a measure of the externalities that an agent generates, and a measure of the degree of supermodularity.

The significance of this result is that by ensuring that agents in a T – set choose strategy 1 the authorities can rule out the inefficient equilibrium: alternatively if the

system is at the inefficient equilibrium then it can be “tipped” to the efficient one by changes of strategy on the part of this subset of players. Such changes may be induced outside the formal structure of the game, for example by side payments or other financial incentives, or alternatively by regulation. As the subset in a smallest T -set may be very small, this may be an attractive policy from a regulator’s perspective. It also provides a possible resolution of the coordination problem for supermodular games, one not previously noted.

Key to our analysis is the effect on agent j ’s payoff of changing strategy from 0 to 1, and how this effect changes when another agent, say i , also changes from 0 to 1. By the increasing differences property (1), we know that the change by i will increase the effect on j ’s payoff of the change by j . Let $s_{-i-j}, 1_i, 0_j$ denote the vector of strategies where all agents k other than i, j are choosing $s_k \in s_{-i-j}$ and i, j are choosing 1 and 0 respectively. Define

$$\Delta_j(i = 0, s_{-i-j}) = u_j(s_{-i-j}, 0_i, 1_j) - u_j(s_{-i-j}, 0_i, 0_j)$$

and

$$\Delta_j(i = 1, s_{-i-j}) = u_j(s_{-i-j}, 1_i, 1_j) - u_j(s_{-i-j}, 1_i, 0_j)$$

These are the returns to agent j from changing from 0 to 1 when agent i chooses either 0 (in the first line) or 1 (in the second line) and everyone else chooses s_{-i-j} .

The difference between these returns is

$$\Delta_{ij}(s_{-i-j}) = \Delta_j(i = 1, s_{-i-j}) - \Delta_j(i = 0, s_{-i-j}) \geq 0 \quad (2)$$

That is,

$$\begin{aligned} \Delta_{ij}(s_{-i-j}) = & [u_j(s_{-i-j}, 1_i, 1_j) - u_j(s_{-i-j}, 1_i, 0_j)] - \\ & [u_j(s_{-i-j}, 0_i, 1_j) - u_j(s_{-i-j}, 0_i, 0_j)] \end{aligned} \quad (3)$$

This is the increase in the return to j 's change of strategy as a result of i 's change of strategy, and from the condition of increasing differences (1) we know that this is positive. It is a measure of the positive externalities generated by a change of i 's strategy, such externalities being guaranteed by increasing differences. As more agents i change their strategy from 0 to 1 there will be a greater increase in utility for the other agents j in the system.

We are interested in the conditions under which an inefficient equilibrium can be tipped to an efficient one, and so focus on equations (2) and (3) when all agents other than i and j are choosing strategy 0 :

$$\begin{aligned} \Delta_{ij}(0) = & [u_j(0^{N-2}, 1_i, 1_j) - u_j(0^{N-2}, 1_i, 0_j)] - \\ & [u_j(0^{N-2}, 0_i, 1_j) - u_j(0^{N-2}, 0_i, 0_j)] \end{aligned} \quad (4)$$

where 0^{N-2} indicates that there are $N - 2$ zeros in the positions other than i and j .

Proposition 1 *Under the assumption 1 of increasing differences there exists a tipping set that tips the equilibrium with all 0s to that with all 1s.*

Proof. Consider the following sequence of inequalities, which link the equilibrium with all zeros to that with all ones in a series of steps in each of which an additional agent changes strategy from zero to one, and which hold by the increasing differences (1) property.

$$\begin{aligned}
 u_i(0^{N-1}, 1_i) - u_i(0^{N-1}, 0_i) &< u_i(0^{N-2}, 1_1, 1_i) - u_i(0^{N-2}, 1_1, 0_i) < \quad (5) \\
 u_i(0^{N-3}, 1_1, 1_2, 1_i) - u_i(0^{N-3}, 1_1, 1_2, 0_i) &< \dots < \\
 u_i(1_1, 1_2, \dots, 1_{N-1}, 1_i) - u_i(1_1, 1_2, \dots, 1_{N-1}, 0_i)
 \end{aligned}$$

If we take the first inequality

$$u_i(0^{N-1}, 1_i) - u_i(0^{N-1}, 0_i) < u_i(0^{N-2}, 1_1, 1_i) - u_i(0^{N-2}, 1_1, 0_i)$$

we see that the payoff to agent i from a strategy change is raised when agent 1 also picks strategy 1. The second inequality

$$u_i(0^{N-2}, 1_1, 1_i) - u_i(0^{N-2}, 1_1, 0_i) < u_i(0^{N-3}, 1_1, 1_2, 1_i) - u_i(0^{N-3}, 1_1, 1_2, 0_i)$$

shows that the payoff to i from the change from 0 to 1 is increased again when agent

2 changes from 0 to 1. The inequalities repeat this process changing one additional agent's strategy each time. The inequalities here hold by the increasing difference property (1).

Note that the first difference is negative

$$u_i(0^{N-1}, 1_i) - u_i(0^{N-1}, 0_i) < 0$$

as the vector of all zeros is a Nash equilibrium so 0 is a best response for i : note also that to the contrary the last difference

$$u_i(1_1, 1_2, \dots, 1_{N-1}, 1_i) - u_i(1_1, 1_2, \dots, 1_{N-1}, 0_i) > 0$$

is positive as the vector of all ones is also a Nash equilibrium and now 1 is a best response. As the sequence of differences starts negative and ends positive it must change sign: let the first positive difference be $u_i(0^{N-t-1}, 1_1, \dots, 1_t, 1_i) - u_i(0^{N-t-1}, 1_1, \dots, 1_t, 0_i)$. Then clearly the first t agents form a T -set. Once they have chosen 1 as a strategy, this is the best response of all other agents. This proves that a T -set exists. ■

In principle we can find the smallest T -set by reviewing this set of inequalities for every possible ordering of agents and finding the ordering for which the change of sign occurs after the smallest number of inequalities. However for a large number of agents this approach could prove extremely time-consuming as the number of

orderings increases exponentially with the number of agents. We can offer more efficient ways of finding the smallest T – set in a special case.

In order to provide a simple characterization of a T – set we focus on the special case in which the difference $\Delta_{ij}(0)$ in equation (4) is independent of the index j , i.e. the effects of i 's change of strategy are symmetric over other agents. In addition we assume that $\Delta_{ij}(s_{-i-j})$ is independent of s_{-i-j} and so does not depend on the strategies chosen by others. These two conditions of symmetry and independence taken together we call assumption A1.

$$\Delta_{ij}(0) = \Delta_{ik}(0) = \Delta_i(0) = \Delta_i \tag{A1}$$

For each agent i , Δ_i is the alteration in the payoff that all other agents get from switching strategy from 0 to 1 as a result of agent i changing from 0 to 1, a uniform externality that i by changing strategy imposes on others when they change strategy.

Given this, agents can be ranked unambiguously by the values of their Δ_i functions, and we assume without loss of generality that they are numbered so that $\Delta_1 \geq \Delta_2 \geq \dots \geq \Delta_N$. An agent's ability to tip the inefficient equilibrium is measured by its Δ , and we show below that the smallest T – set consists of agents with the greatest Δ s.

Let $\{0^t, 1^{N-t-1}, 1_k\}$ denote the following vector: the k – th coordinate is 1, t coordinates are zero, all others (of which there are $N - t - 1$) are 1, and the first $N - t - 1$ coordinates are 1 if $k > N - t$ and the first $N - t$ are 1 otherwise.

From (2) and (3) and (A1) we can write

$$u_j(0^{N-K-1}, 1^K, 1_j) - u_j(0^{N-K-1}, 1^K, 0) = u_j(0^{N-1}, 1) - u_j(0^{N-1}, 0) + \sum_{i=1}^{i=K-1} \Delta_i \quad (6)$$

Hence finding a t such that $u_i(0^{N-t}, 1_1, \dots, 1_t, 1_i) - u_i(0^{N-t}, 1_1, \dots, 1_t, 0_i) > 0$ is the same as finding a t such that $u_j(0^{N-1}, 1) - u_j(0^{N-1}, 0) + \sum_{i=1}^{i=t-1} \Delta_i > 0$ or $\sum_{i=1}^{i=t-1} \Delta_i > u_j(0^{N-1}, 0) - u_j(0^{N-1}, 1)$. From this we can readily prove:

Proposition 2 *Given A1, if a smallest T - set exists then for some integer F it consists of the first F agents when agents are ranked by the value of Δ_i .*

Proof. If $F < N$ is a T - set then for all $j > F$ we must have

$$u_j(0^{N-F-1}, 1^F, 1_j) - u_j(0^{N-F}, 1^F) \geq 0$$

By (6) above this is equivalent to

$$\sum_{i=1}^{i=F-1} \Delta_i \geq u_j(0^{N-1}, 0) - u_j(0^{N-1}, 1) \quad \forall j > F \quad (7)$$

To construct the smallest T - set we need to find the smallest number F for which (7) holds. Clearly we get this by ranking agents by the size of Δ_i and choosing first those with the largest value of Δ_i , i.e. those that create the largest externalities or that exhibit increasing differences to the greatest degree. ■

Proposition 2 shows that the agents that are most capable of changing the game's

equilibrium are those that generate the largest externalities to others, and that the ability to change the equilibrium depends on the game being sufficiently supermodular or on the degree of increasing differences being great enough. The terms Δ_i are measures of the degree of increasing difference, and assumption *A1* places a structure on these so that they are symmetric across agents. This structure is necessary for the simplicity of our arguments but not for the basic intuition that increasing differences contribute to tipping, as Proposition 1 shows. Within the structure defined by *A1* we can say that increasing differences being sufficiently large in the sense of (7) is necessary and sufficient for tipping of the inefficient equilibrium. A numerical example of tipping at the inefficient equilibrium of a super modular game is given in [6].

It is possible to establish results like Proposition 2, but more complex ones, with weaker assumptions than *A1*. Suppose for example that we drop the independence assumption, namely that $\Delta_{ij}(s_{-i-j})$ is independent of s_{-i-j} . In this case in stating and deriving a proposition analogous to Proposition 2 we need to reorder the agents by the size of Δ_i after each selection of a member of the tipping set, as the change of strategy from 0 to 1 by one agent can alter the ranking of the agents still choosing 0 by their $\Delta_i s$. In forming the tipping set at each stage we add the remaining agent whose Δ_i is greatest given the strategies now in place by all other agents, and this gives a more general but less parsimonious version of Proposition 2. If we drop the symmetry condition we are back with the general case of Proposition 1.

3 Cascading

A cascade is a sequence of events in which a change of strategy by one agent leads another to change, the changes of the two together lead a third to change, and so on. It is a version of the classic domino effect. Avinash Dixit models this well and we follow his framework [3]. In our context a cascade will begin from an equilibrium where all agents choose $s_i = 0$. A cascade of length k is a situation where:

- if 1 were to change from 0 to 1 but all others were to remain at 0 then 2's best response would be 1
- if 1 and 2 were to choose 1 and all others 0, then 3's best response would be 1.
- if 1, 2 and 3 were to choose 1 and all others 0, then 4's best response would be 1
- and so on up to agent k . The strategy tuple in which agents 1 through k choose 1 and all others choose 0 is a Nash equilibrium.

If we think of the game as one in which moves are made sequentially by players in ascending order, we will see that if the first mover chooses 1 (perhaps as a result of factors outside the game as we have defined it, such as policy inducements) then the second mover chooses 1 and so on up to and including the k -th mover and thereafter all will choose 0 and the outcome will be an equilibrium. If the only equilibria involve either all zeros or all ones then the outcome of such a cascade will be the equilibrium

with all 1s, and this will be attained from the equilibrium of zeros by persuading agent number one to change strategy.

Formally we have a cascade of length k at the Nash equilibrium $\{0, 0, \dots, 0\}$ if agents can be numbered so that agent 2's best response to $\{1, s_2, 0, \dots, 0\}$ is $s_2 = 1$, agent 3's best response to $\{1, 1, s_3, 0, \dots, 0\}$ is $s_3 = 1$, and for all agents j for $j \leq k$ the best response to $\{1, 1, \dots, s_j, 0, \dots, 0\}$ is $s_j = 1$, and in addition $\{1, 1, \dots, s_k = 1, 0, \dots, 0\}$ is a Nash equilibrium. Using the framework and assumptions of the previous section we can set out sufficient conditions for a cascade of length k .

We can give a formal characterization of the conditions for a cascade of length k as follows:

Proposition 3 *A cascade of length k occurs if*

$$\Delta_{j-1} \geq u_j \left(\overbrace{1, \dots, 1}^{j-2}, 0, \dots, 0 \right) - u_j \left(\overbrace{1, \dots, 1}^{j-2}, 0, 1, 0, \dots, 0 \right)$$

for all $j \leq k$.

Proof. For a change by agent 1 to change agent 2's strategy we need that

$$u_2(1, 1, 0, \dots, 0) - u_2(1, 0, \dots, 0) = u_2(0, 1, 0, \dots, 0) - u_2(0, \dots, 0) + \Delta_1 > 0$$

or

$$\Delta_1 > u_2(0, \dots, 0) - u_2(0, 1, 0, \dots, 0)$$

Similarly for a change by agent 2 to change 3's strategy

$$u_3(1, 1, 1, 0..0) - u_3(1, 1, 0..0) = u_3(1, 0, 1, 0..0) - u_3(1, 0, ..0) + \Delta_2 > 0$$

or

$$\Delta_2 > u_3(1, 0, ..0) - u_3(1, 0, 1, 0..0)$$

The proposition follows by repeating this argument. ■

Cascading, like tipping, depends on a game exhibiting “enough increasing difference.” A numerical example of cascading from the inefficient to the efficient equilibrium of a supermodular game is given in [6].

4 Schelling's work

Schelling [13] provides a number of examples of the role of a critical mass in inducing tipping: in these examples individuals make a decision about being part of process or group based on what they see others doing. A key example is given by individuals' decisions about whether to reside in a neighborhood, which they do if there are enough others like themselves who are already there. Schelling's most famous example, of racial segregation in residential districts, was essentially dynamic, with a sequence of changes evolving over time. However it is possible to capture much of what was interesting in and essential to that model with a static formulation identical to that used above.

Consider a population of P people of a certain type living in a neighborhood. Each has two possible strategies - stay S or move M . The payoff to each depends on how many others of the group do the same: the payoff to staying is the number of others who stay, $\#(S)$, and the payoff to moving is likewise the number of others who do this, $\#(M)$. Clearly all staying or all moving are both Nash equilibria, and if $\#(M) > \#(S)$ then the best response of all who have not moved is to move, so that we have the possibility of tipping. This game displays increasing differences, as the payoff to changing from S to M increases with the number of people who have already changed.

5 Airline security

We now illustrate our results on tipping and supermodularity in a case that we and others have studied extensively, that of airline security (see Howard Kunreuther and Geoffrey Heal [9], Heal and Kunreuther [6] and Michael Kearns [8]). Each airline is concerned with how large an investment it should incur to reduce the likelihood of a dangerous bag or passenger being on board one of its airplanes. Each knows that even if it invests in security screening there is still a chance that a dangerous bag or passenger could transfer to it from another airline with lax inspection procedures. There is nothing it can do to stop this process short of inspecting all passengers and bags transferred from other airlines, a time consuming and costly process followed only by El Al.. The Pan Am 103 crash illustrates this case well: the bag that caused

the accident was loaded at Gozo Airport in Malta, with lax security, transferred to a Pan Am feeder in Frankfurt and then to Pan Am 103 in Heathrow, set to explode at over 28,000 feet. There was nothing that Pan Am could have done to prevent this disaster short of inspecting all bags transferred to its planes from other airlines.

5.0.1 The Model

Kunreuther and Heal (2003) [9] have developed a game theoretic model to show that each airline will have less incentive to invest in security if it knows that other airlines have not invested. In the context of the above model, can one specify conditions for tipping an equilibrium from one where none invest to an equilibrium where all invest in security?

There are $n \geq 2$ separate airlines. During the course of a given time period each airline makes a certain number of trips, each of which is identical. Consider a given plane trip initiated by airline i , and assume that the airline has made no investments in security systems. Let p_{ij} be the probability that on any trip a bag containing a bomb is loaded onto airline i and is then transferred to airline j , exploding on j . If $i = j$, then p_{ii} represents the probability that an airline loads a bag with a bomb and this explodes on its own plane. Let $p_i = \sum_j p_{ij}$ and $\tilde{p}_i = \sum_{j \neq i} p_{ij}$. Thus p_i is the probability of airline i loading a bomb that explodes on its own plane and \tilde{p}_i is the probability that it loads a bomb that explodes on another airline - a measure of the risk that it poses to others. We expect that $p_i < 1$ so that there is some chance that

the airline does not load a bag with a bomb that explodes. Each airline can either invest in a security system (strategy = S) at a cost per trip of $c_i > 0$ or not invest N . Security systems are assumed to be completely effective so that they eliminate the chance of a bomb coming through the airline's own facility. In the event that a bomb explodes on a plane, the loss is $L > 0$. The initial income of an airline is $Y > c_i \forall i$.

In the case of just airlines A_1 and A_2 that exchange passengers and baggage maximizing expected profits, this framework gives rise to the following payoff matrix showing the outcomes for the four possible combinations of N and S . If both airlines invest in security systems then their payoffs per trip are just their initial incomes net of the investment costs. If A_1 invests and A_2 does not, then A_1 has a payoff of income Y minus investment cost c_1 minus the expected loss from a bomb transferred from A_2 that explodes on A_1 (i.e., $p_{21}L$), while A_2 has a payoff of income Y minus the expected loss from a bomb loaded and exploding on 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	$Y - p_{11}L, Y - c_2 - p_{12}L$	$Y - p_{11}L - (1 - p_{11})p_{21}L, Y - p_{22}L - (1 - p_{22})p_{12}L$

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

$$c_1 < p_{11}L \text{ and } c_1 < p_{11} [1 - p_{21}] L \quad (8)$$

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.

Following the model developed in section 2, we need to identify the change in the return to airline 1 when it invests in security as a result of investment by airline 2. From the payoff matrix it is easy to see that this is

$$\Delta_{21} = p_{11}p_{21}L$$

By similar reasoning we can show that for the cases of three and four airlines when all are not investing and airline 2 changes from not investing to investing the expressions for the change in 1's payoff are respectively

$$\Delta_{21} = p_{11}p_{21}L (1 - p_{31}), \quad \Delta_{21} = p_{11}p_{21}L (1 - p_{31}) (1 - p_{41})$$

For the general case of N agents the change in the return to j investing as a result

of i investing when no other agents are investing is

$$\Delta_{ij} = p_{jj}p_{ij}L \prod_{k \neq i,j} (1 - p_{ik}) \quad (9)$$

If agents other than i, j are investing, say agents in $\{S\}$, then they are excluded from the product:

$$\Delta_{ij} = p_{jj}p_{ij}L \prod_{k \neq i,j, j \notin \{S\}} (1 - p_{ik})$$

However, we are interested in tipping an equilibrium at which no agents are investing, so we are interested in the case in which S is empty and the first version of the formula applies. The expression (9) has a natural intuitive interpretation. The term $p_{ij}L \prod_{k \neq i,j} (1 - p_{ik})$ is the expected cost to agent j of a security failure at agent i , conditional on there not being a security failure at another agent $k \neq i, j$. The higher this expected cost the greater the increase in j 's expected gains from investing if i invests in protection as well. Multiplying this expected gain by p_{jj} normalizes this value given the chance of a loss because of a security failure at j and determines the increase in expected profit to j from i investing in security. Consider the extreme case where $p_{jj} = 0$. For the two agent case if $p_{11} = 0$ then, as is clear from the payoff matrix and expression (8), there is no return to agent 1 to investing in security on its own. It would then never be optimal for agent 1 to invest in security whether or not agent 2 invested. More generally the smaller the value of p_{jj} , the less reason agent j will have to invest in security on its own and hence the less likely that this agent will

be a candidate for being tipped into investing should other agents such as i invest in security (9).

With this characterization of the Δ_s we are in a position to apply the results of the previous sections to understand the tipping and cascading possibilities for the airline security problem. Assume, following Assumption A1, that Δ_{ij} is independent of j so that $\Delta_{ij} = \Delta_i, p_{jj} = p$ and $p_{ij} = q_i$. In this case 9 becomes:

$$\Delta_{ij} = \Delta_i = pq_i L \prod_{k \neq i, j} (1 - p_k) \quad (10)$$

If there are two Nash equilibria with either everyone or no one investing in protection, then by Proposition 1 the smallest T – set consists of the first K agents when agents are ranked by the value Δ_i . Agents will have higher values of Δ_i if they have high values of q_i .

5.1 Applications to Other IDS Problems

As shown in Kunreuther and Heal [9] there are a wide range of problems that exhibit features of supermodularity where tipping could occur. One area that naturally falls into this class is computer security. Here the central issue is the incentive each agent has to invest in protecting itself against a possible virus, when it knows that it may be infected from other agents. The following example adapted from Kearns [8] illustrates this problem. Imagine the user population of a large organization in which each individual has a desktop computer with its own local software and memory, but

in which parties also maintain important data files or documents on a shared disk drive accessible to the entire organization.

From the perspective of the organization, the primary security concern is that an intrusion (whether by a piece of malicious software or a human hacker) might erase the contents of this shared hard drive. Each user's desktop computer and its contents—including E-mail, downloaded programs or files, and so on—is a potential point of entry for such intrusion. Each user must implicitly decide about many aspects of their individual security practices: how often they change their password (and how secure those passwords are against dictionary and other common attacks), whether they enable encryption in their web and e-mail communications, their care in not downloading suspicious files and programs, their anti-virus software maintenance, and many other features, each of which takes time and hence is costly. The vulnerability of the shared hard drive is determined by the collective behavior along these dimensions. Hence if an agent invests in rigorous security precautions, her investments can be compromised by a failure to do likewise on the part of just one other. As it may be the case that some other agents do not store valuable data on the shared drive, their incentives to adopt best-practice security measures may be small. So we again have an interdependent security problem, with supermodularity of the associated game and the possibility of tipping from an equilibrium where none take security seriously to one where all do.

A related problem is one where each division in a decentralized firm needs to

determine whether it wants to incur the costs of investing in risk-reducing measures knowing that other divisions in the same firm may not follow suit.(see [10] for more details). By not investing in protection there is some probability p_i that division i will have a large loss that could cause its division to fail and $q_i \leq p_i$ is the probability that the loss would be so large that the entire firm would become insolvent. Two recent examples are Nick Leeson operating in the Singapore futures market division causing the collapse of Baring's Bank and Arthur Andersen being brought into bankruptcy by the actions of its Houston branch. The losses L in this case are the costs that managers and other employees of the division will incur if their division goes bankrupt. These include the search costs for new employment and other negative features associated with losing ones job including loss of reputation. The ranking of agents in forming a $T - set$ is similar to that in proposition 2 so that one would want to find ways to encourage those divisions in the firm who create the largest expected negative externalities to be the first ones to invest in protection and thus induce other divisions to follow suit.

References

- [1] 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.

- [2] Crawford, Vincent and Hans Haller, Hans 1990. "Learning How to Cooperate: Optimal Play in Repeated Coordination Games" *Econometrica* Vol 58: Issue 3 (May 1990) 571 - 595.
- [3] Dixit Avinash K 2002. "Clubs with entrapment." Available at www.princeton.edu/~dixitak/home. Published in *American Economic Review* Vol.. 93 No. 5 pp. 1824-1829.
- [4] Easterly, William 2004 "Empirics of Strategic Interdependence" February (mimeo)
- [5] Gladwell Malcolm 2000. *The Tipping Point* Little Brown and Co.
- [6] Heal Geoffrey and Kunreuther Howard.(2005). "IDS Models of Airline Security" *Journal of Conflict Resolution* 41:201-17.
- [7] Heller, Walter 1986 "Coordination Failure in Complete Markets with Applications to Effective Demand". In *Equilibrium Analysis: Essays in Honor of Kenneth J. Arrow Vol II*, ed. W.P.Heller, R.M. Starr and D.A. Starrett, Cambridge University Press., 1986.
- [8] Kearns, Michael 2005. "Economics, Computer Science, and Policy". *Issues in Science and Technology*, Winter pages 37-47.

- [9] Kunreuther Howard and Geoffrey Heal 2003. "Interdependent Security." *Journal of Risk and Uncertainty*, Special Issue on Terrorist Risks, Vol. 26 No. 2/3 (March/May): 231-249.
- [10] Kunreuther Howard and Geoffrey Heal 2005 "Interdependencies within an Organization" in B. Hutter and M. Powers (ed.) *Organizational Encounters with Risk* (Cambridge University Press)
- [11] Milgrom Paul and John Roberts. 1990. "Rationalizability, Learning and Equilibrium in Games with Strategic Complementarities" *Econometrica* 58: 1255-77.
- [12] Schelling, Thomas 1971. "Dynamic Models of Segregation" *Journal of Mathematical Sociology* 1:143-86.
- [13] Schelling, Thomas 1978. *Micromotives and Macrobehavior*. New York: Norton
- [14] Topkis, Donald 1978 "Minimizing a Submodular Function on a Lattice" *Operations Research* 26: 305-21.
- [15] Topkis, Donald 1979 "Equilibrium Points in Nonzero-Sum n-Person Submodular Games" *Siam Journal of Control and Optimization* 17:773-87.
- [16] Vives, Xavier. *Journal of Economic Literature* Vol.. XLIII (June 2005)