Risk Management Strategies for Dealing with Interdependencies[1]

Howard Kunreuther[2]

Paper presented at Conference on
Network-based Strategies and Competencies
Wharton School
University of Pennsylvania
November 8-9, 2007

Working Paper # 2007-11-20

3730 Walnut Street, Jon M. Huntsman Hall, Suite 500
Philadelphia, PA 19104
USA
Phone: 215-898-4589
Fax: 215-573-2130
http://opim.wharton.upenn.edu/risk/

[1] This paper is based on several papers written jointly with Geoffrey Heal as well material that was distributed at a conference on *Interdependent Security (IDS: Theory and Practice)* co-organized with David Bell and Vicki Bier and sponsored by CREATE (USC), the Center for Human Performance and Risk Analysis (U. of Wisconsin) and the Wharton Risk Management and Decision Processes Center. Support from NSF Grant No. CMS-0527598 and the Wharton Risk Management and Decision Processes Center is gratefully acknowledged.

[2] Cecilia Yen Koo Professor of Decisions Sciences and Public Policy at the Wharton School, University of Pennsylvania and co-director of the Wharton Risk Management and Decision Processes Center.
E-mail address: Kunreuther@wharton.upenn.edu

CITATION AND REPRODUCTION

This document appears as a Working Paper of the Wharton Risk Management and Decision Processes Center, The Wharton School of the University of Pennsylvania. Comments are welcome and may be directed to the authors.

This paper may be cited as: Howard Kunreuther, "**Risk Management Strategies for Dealing with Interdependencies**," Risk Management and Decision Processes Center, Wharton School of the University of Pennsylvania, November 2007.

THE WHARTON RISK MANAGEMENT AND DECISION PROCESSES CENTER

Since its creation 25 years ago, the mission of the Wharton Risk Management and Decision Processes Center has been to carry out a program of basic and applied research to promote effective corporate and public policies for low-probability events with potentially catastrophic consequences. The Risk Center has focused on natural and technological hazards through the integration of risk assessment, risk perception and risk financing with risk management strategies. After 9/11, research activities extended also to national and international security issues (e.g., terrorism risk insurance markets, protection of critical infrastructure, global security).

Building on the disciplines of economics, decision sciences, finance, insurance, marketing and psychology, the Center's research program has been oriented around descriptive and prescriptive analyses. Descriptive research focuses on how individuals and organizations interact and make decisions regarding the management of risk under existing institutional arrangements. Prescriptive analyses propose ways that individuals and organizations, both private and governmental, can make better decisions regarding risk. The Center supports and undertakes field and experimental studies of risk and uncertainty to better understand the linkage between descriptive and prescriptive approaches under various regulatory and market conditions. Risk Center research investigates the effectiveness of strategies such as risk communication, information sharing, incentive systems, insurance and regulation.

The Center is also concerned with training decision makers and actively engaging multiple viewpoints, including the expertise of top-level representatives in the world from industry, government, international organizations, interest groups and academics through its research and policy publications and through sponsored seminars, roundtables and forums.

More information is available at http://opim.wharton.upenn.edu/risk

Risk Management Strategies for Dealing with Interdependencies

1. Introduction

In an interdependent world, the risks faced by any one agent depend not only on that agent's own choices but also on those of others. More specifically, the economic incentive of any agent to invest in protection depends on how she expects the others to behave. The strategies can be risk-reducing measures as well as information-gathering and preparedness activities. The fact that such events are typically probabilistic, and that the risk that one agent faces is often determined in part by the behavior of others, gives a unique and complex structure to the incentives that agents face to reduce their exposures to these risks that come under the heading of interdependent security (IDS).

For many IDS problems, if an agent thinks that others will *not* invest in protection, then this reduces the incentive for her to do so. On the other hand, should she believe that others will invest in security, it may be best for her to do so also. So, it is often the case that there are two equilibria: (1) no one invests in protection, even though all would be better off if they had incurred this cost or (2) everyone invests in protection. This situation does not have the structure of a prisoners' dilemma game, though it has many similarities.

This paper characterizes the nature of the interdependency problem and suggests risk management strategies for improving both individual and social outcomes. The next section of the paper outlines a series of IDS scenarios to illustrate the range of problems that fall under this rubric. Section 2 focuses on the problem of a firm with more than one decentralized division and develops a simple game theoretic model to show how the expected profits of each of the individual divisions would be improved had all of them invested in risk-reducing measures. We then examine how one might induce tipping or cascading by either subsidizing or fining one of the divisions so that it has an economic incentive to invest in protection, leading the others to follow suit. Section 3 introduces risk management strategies to improve individual and social welfare. In particular, I focus on coordination measures within a decentralized firm (e.g., creating corporate culture) as well as within an industry (e.g., private trade associations) to induce cooperative behavior. The public sector can also play an important role in this regard through interventions such as taxes, subsidies, insurance and regulations to deal with the negative externalities caused by interdependent security. The paper concludes with suggestions for future field and experimental research in this area.

2. IDS Scenarios

In some of the IDS scenarios below, weak links in the system may lead to suboptimal behavior by everyone. In such cases there may be ways of inducing tipping and cascading so that everyone's welfare is improved. One may then want to determine the nature of critical coalitions that can tip the entire system. Is there any agent (e.g., firm, individual) or group of agents one should focus attention on? More generally, what types of private sector coordination measures (e.g., private trade associations) and public sector interventions such as taxes, subsidies, fines, regulations and well-enforced standards are appropriate for dealing with the negative externalities caused by interdependent security.

Example 1: Encouraging installation of theft alarms[3]

A person who is considering investing in a car alarm system needs to balance the cost of installing the alarm and maintaining it against the benefits associated with deterring the thief. He has the option of investing in a silent alarm system (e.g., LoJack, where a small radio frequency transceiver installed in the car informs the police that a car has been broken into so they can track it) and does not publicly announce that he has done so. Over time, thieves would learn that X percent of the cars were alarmed with a LoJack-type system. By investing in a silent alarm, others in the community benefit: If enough people take this action then thieves will look for areas where they have a lower probability of being caught. On the other hand, the car owner could install an alarm system and announce it publicly. The expected benefits to an individual of alerting others that she has a theft protection system could be much greater than a secret alarm, because the thief knows that he should attack another vehicle. He does not have this information if a silent alarm is installed.

Example 2: Protection of shared network resources[4]

In many workplaces, there exists a complex network of shared resources (such as files, disks, peripheral devices, and bandwidth) along with individual resources (such as desktop machines). The vulnerability of the shared resources to various security exploits often depends strongly on the collective actions used to protect individual resources. For example, a shared disk may be erased by a virus entering the local network through the desktop machine of a user who failed to update his or her anti-virus software signatures. An interesting example is one in which the shared resource is bandwidth, and where users whose machines are infected with a variety of "malware" can surreptitiously consume huge amounts of bandwidth, at the expense of all other users of the system. Such problems are common among residential commercial cable subscribers to Internet access**.**

Example 3: Baggage transfer security[5]

An airline has to determine whether it wants to invest in baggage security, knowing that even if it takes this action, it may face a security risk from a dangerous bag loaded onto its plane by another airline. It faces this risk unless it inspects all transferred bags. Lest this point be considered only theoretical, recall the crash of Pan Am 103 in 1988. Terrorists checked a bag containing a bomb in Malta on Malta Airlines, which had minimal security procedures. The bag was transferred in Frankfurt to a Pan Am feeder

---

[3] See Ayers and Levitt (1998) for more details on LoJack.
[4] See Heal et al., (2006) and Kearns (2005) for more details on this scenario
[5] See Heal and Kunreuther (2005a) for more details on this scenario

line, and then loaded onto Pan Am 103 in London's Heathrow Airport. The bomb was designed to explode above 28,000 feet, a height normally first attained on this route over the Atlantic Ocean. Thus, in this case, the terrorists deliberately exploited the widely varying security procedures across the airlines. This problem is common to other transportation modes, where there are interconnections between nodes in the network.

Example 4: Reducing the risk of power failures[6]

Consider a utility that is part of is part of an integrated system, (i.e. the power grid) and wants to determine whether to invest in additional capacity or security measures (e.g., trimming vegetation near distribution lines) to reduce the chance of a power failure. In any highly interdependent system, such as the power grid, there is a systemic tendency to under-invest in reliability. A consequence of the interdependency is that a part of the cost of a failure, perhaps a large part, is passed on to competitors and their customers. In the case of the August 2003, power failures in the northeastern U.S. and Canada, the costs of a failure at an Ohio utility, were passed on to other utilities and customers over the northeastern US and southeastern Canada. There are two routes to a solution. One, the property rights approach, is to hold a utility responsible for the full costs of a service failure, wherever it occurs. The second, the regulatory approach, is to mandate minimum reliability standards with monitoring and serious penalties for non-compliance. By forcing each utility to bear the full costs of its shortcomings, the first route provides a clear incentive to avoid failures. The second seeks to prevent them through regulatory action.

Example 5: Supply-chain management[7]

The effects of supply-chain disruptions (whether due to natural disasters, terrorists, or other unexpected events) on the profitability of supply-chain participants are now recognized as being potentially catastrophic. Shipping delays and other supply-chain disruptions during the 1990s showed that companies experiencing such disruptions underperformed their peers significantly in stock performance, as well as in operating performance (as reflected in costs, sales, and profits). Coping with the management challenges of such disruptions and weak links in the supply chain is, however, a difficult matter, as the interdependencies require cooperative activity and monitoring across the supply chain in ways that are not captured in the traditional metrics of price, time/responsiveness, and product quality. For example, firms who are protecting themselves against disruptions due to the avian flu may find that their own measures may be of little use unless accompanied by protections on the part of their suppliers, the suppliers' suppliers, etc. Another example of interdependency in the supply chain is protecting food or agricultural products against physical contamination (bacteria, toxins, etc.) such as protecting milk from botulinum toxin, and protecting cattle from foot-and-mouth disease. In these cases, the product could be distributed widely and poison thousands.

---

[6] SeeFeinstein (2006) for more details on this scenario
[7] See Heal et al., (2006) for more details on this scenario

Example 6: Vaccination against diseases[8]

If people face the possibility of catching an infectious disease and have the option of a vaccine, how many will choose to be vaccinated? Intuition suggests that not all will. If most people are vaccinated, there is little incentive for the unvaccinated individuals to join them, if they can only catch the disease from another human. In this case, if everyone else is vaccinated, then it is optimal for the last person not to be vaccinated, since she faces no risk and can free-ride on the herd immunity of the community. Likewise, if no one is vaccinated, everyone has a strong incentive to be vaccinated if there is a risk of exposure. In this case, no one being vaccinated is an equilibrium solution only if the cost of vaccination is extremely high. This seems to suggest that there should be an interior solution with some but not all of the population being vaccinated. Who, and how many, should be vaccinated? Can one identify individuals who are more susceptible to the disease and/or are more likely to spread it to others who are prime candidates for protection if there is a limited supply of the vaccine?

Example 7: Meltdown of a nuclear reactor[9]

Assume that each country has one nuclear reactor, and that if it invests in a set of safeguards, the chances of an accident from the power plant are reduced to zero. We imagine a group of small adjacent countries (e.g., Belgium, Holland, and Luxembourg, or Latvia, Lithuania, and Estonia), where a meltdown in any one will lead to radioactive contamination in all of them. What role could international compacts or trade associations play to ensure that all countries invest in safeguards? A related issue is the externalities from an accident that does not result in contamination. For example, in the U.S., an accident at any one plant is likely to lead to costly regulatory interventions at all plants. That is why the Institute of Nuclear Power Operations was founded after the incident at Three Mile Island—to serve as a self-policing arm for the industry, so that well-performing plants would not be held hostage to the safety problems of plants with poorer safety records.

Example 8: Environmental treaties[10]

Suppose that countries are asked to sign a treaty to reduce some environmental risk, such as global warming or atmospheric pollution. There is a net cost to any one country for adopting the treaty, but potential benefits to the planet if enough countries take this action. What incentive is there for any one country to adopt the treaty if it knows that a number of other countries will not join? How can one convince countries with leverage to sign the treaty to induce others to follow suit? There are equity-efficiency tradeoffs that may have to be addressed here. For example, one can envision that it might be economically more efficient for only a subset of countries to take preventive actions by being part of a treaty, but more equitable and politically saleable for all countries to sign the treaty.

---

[8] See Heal and Kunreuther, (2002) for more details on this scenario
[9] See Heal and Kunreuther, (2005b) for more details on this scenario
[10] See Barrett (2003) for more details on this scenario

Example 9: Interdependent critical infrastructures[11]

In the wake of 9/11 and Hurricane Katrina, the private and public sector share an interest in making social and economic systems less vulnerable to disasters. There is a growing interest in protecting critical infrastructure that assure the social and economic continuity of the nation (transportation, water distribution, telecommunication, electricity, emergency services, finance sector, etc.). One of these challenges is the existence of interdependent operations between multiple infrastructures in different sectors. For example, financial systems and emergency services are highly dependent on telecommunication operations, which are highly dependent on electricity. When the interdependencies cut across sectors, the nature of the risks are often not well understood so that they pose special policy challenges.

Example 10: Reducing the Risk of Wildfire[12]

Suppose you have built a home in one of the areas outside of San Diego that is designated as a red or high hazard zone with respect to wildfires. You are aware of the potential for fire and decide to install a tile or metal roof to reduce the chances that a fire will damage or destroy your house. Unfortunately, your next door neighbor, who has not thought about the possibility of fire because it "won't happen to me" installs a shake roof made of cedar. This roof is pleasing to the eye and less expensive than either a tile or metal roof. However, in a dry desert climate, the roof is like a match stick that could be ignited by a spark and cause the house to burn. Such a disaster will not be confined to your neighbor's house but will very likely to spread to yours and others even though you have a tile or metal roof.

Example 11: Protecting a firm against catastrophic losses[13]

Consider the possible bankruptcy of a firm due to a catastrophic loss to one of its plants or divisions. An example would be a Bhopal-like accident to a chemical plant where the losses are so large that they cause bankruptcy of the entire operation of the firm nationwide, or even worldwide. Another example is an ownership group such as Lloyd's, which controls a number of syndicates, all operating in a semi-autonomous fashion. If one of the syndicates experiences a severe enough loss, it can lead the ownership group to declare bankruptcy. Recently, Arthur Andersen was sent into bankruptcy by the actions of its Houston branch. Similarly, several years ago, Barings was destroyed by the actions of a single trader in its Singapore division. Given such an institutional structure, what economic incentive does any unit have to incur the costs of protective measures that adversely affect its balance sheet, if other divisions in the organization are not taking similar actions? A culture of risk-taking can spread through the firm, because knowledge that a few groups are taking large risks reduces the incentives that others have to manage their operations carefully.

---

[11] See Auserwald et al ,.(2006) for more details on this scenario

[12] See Spyratos, V., P. Bourgeron, and M. Ghil, (2007) for more details on this scenario.

[13] See Kunreuther and Heal (2005) for more details on this scenario

3. Characterizing the Problem—Investing in a Chemical Plant[14]

I now elaborate on the issues of interdependency characterized in Example 11 by utilizing chemical plant protection as an illustrative example by first focusing on a two-division example and then extending it to multiple divisions.

The Two Division Problem

Consider two identical independently operating divisions in the BeSafe chemical firm: A_1 and A_2, each maximizing its own expected returns and having to choose whether to invest in a protective measure. Such an investment would reduce the probability of a catastrophic chemical accident to one of its plants. Let p be the probability of an accident that bankrupts the division and $q \le p$ be the probability that the accident bankrupts the entire firm. Note that q and p are not independent of each other. The loss from an accident to the participants in the division is L. One should view L as the costs that managers and other employees of the division will incur if their division goes bankrupt. These include the costs to search for new employment and other negative features associated with losing one's job, including loss of reputation.[15] The expected loss from such a catastrophic accident to the participants in the division is thus pL. If the division has invested in protective measures at an upfront cost of c, then the chances of this accident is reduced to zero.[16]

Suppose Division 1 has invested in protection. There is still an additional risk q that BeSafe will go bankrupt if Division 2 has not taken this precautionary measure. In other words, the employees in Division 1 may lose their jobs because of the carelessness of Division 2. In this sense, Division 2 can contaminate other parts of the organization by **not** protecting its plants against a catastrophic accident. Similarly, Division 1 can contaminate Division 2 if it fails to adopt adequate protection.

Let Y be the assets of each division before any expenditure on protection or any losses during the year from the risks faced. Employees in each division are assumed to receive bonuses at the end of the year that are proportional to the size of the division's ending assets. The cost of investing in security is $c<Y$. Assume that each division has two choices: invest in security, **S**, or do not invest, **N.** We can construct a simple 2x2 matrix illustrating what happens to the expected returns of each division as a function of the choices each one makes. The four possible paired outcomes are shown in Table 1:

[14] This section draws heavily from Kunreuther and Heal (2005)

[15] See Greenwald and Stiglitz (1990) for a further discussion of the costs of bankruptcy that professional managers in firms suffer should the firm go bankrupt.

[16] We have assumed the risk to be zero to simplify the exposition. The qualitative results do not change if there is still a positive probability of an accident with a loss of *L* after precautionary measures have been adopted.

Table 1: Expected Returns Associated with Investing and Not Investing in Security

		Division 2 (A_2 ***)***	
		S	***N***
Division 1 (A_1 ***)***	***S***	*Y-c, Y-c*	*Y-c-qL, Y -pL*
	N	*Y-pL, Y-c-qL*	*Y –[pL + (1-p) qL], Y –[pL+ (1-p)qL]*

To illustrate the nature of the expected returns, consider the upper left-hand box where both divisions invest in security (**S,S)**. Then each division incurs a cost of *c* and faces no possible catastrophic accidents so that each of their net returns are *Y-c*.

If A_1 invests and A_2 does not then this outcome is captured in the upper right hand box (**S,N**). Here A_1 incurs an investment cost of *c* but there is still a chance *q* that a catastrophic accident will occur in A_2 causing BeSafe to go bankrupt so that A_1's expected loss from damage originating elsewhere is *qL*. This type of contamination imposed by A_2 on A_1 is referred to in economics as a ***negative externality.*** A_2 incurs no cost of protecting itself and faces no risk of bankruptcy from A_1, but does face the risk of damage originating in one of its plants with an expected loss of *pL*. The lower left box (**N, S)** has payoffs which are just the mirror image of these.

Suppose that neither division invests in protection (**N, N) –** the lower right-hand box in Table 1**.** Then each has an expected return of *Y- pL -(1-p)qL*. The expected losses can be characterized in the following manner. The term *pL* reflects the expected loss originating from an accident in ones own division. The second term reflects the expected loss from an accident originating at the other division that bankrupts the firm (*qL)* and is multiplied by *(1-p)* to reflect the assumption that bankruptcy to a division can only occur once. In other words, the risk of contamination only matters to a division when that division does not have a catastrophic accident originating at home.

Since each division wants to maximize the expected returns to its own employees, the conditions for it to invest in protection against a catastrophic accident are that $c<pL$ and $c<p(1-q)L$. The first constraint is exactly what one would expect if BeSafe consisted of a single division: that is, the cost of investing in protection must be less than the expected cost to its employees from a catastrophic accident. Adding a second division tightens the constraint by reflecting the possibility of contamination from others. This possibility reduces the incentive to invest in protection. Why? Because in isolation, investment in protection buys the employees in the division freedom from bankruptcy. With the possibility of contamination from others it does not. Even after investment there remains a risk of bankruptcy emanating from the other division. Investing in protection buys you less when there is the possibility of contamination from others.

This solution concept is illustrated below with a numerical example. Suppose that *p= .1, q=.05, L=1000, c= 98.* The matrix in Table 1 is now represented as Table 2.

Table 2: Expected Returns Associated with Investing and Not Investing in Security
Illustrative Example: *p= .1 q=.05, L=1000* and *c= 98*

		Division 2 (A_2)	
		S	***N***
Division 1 (A_1)	***S***	*Y-98, Y-98*	*Y-148, Y -100*
	N	*Y-100, Y-148*	*Y-145, Y- 145*

One can see that if A_2 has protection (**S**), then it is worthwhile for A_1 to also invest in security since its expected losses will be reduced by $pL= 100$ and it will spend *98* on the security measure. However, if A_2 does not invest in security (**N**), then there is still a chance that A_1 incurs a loss. Hence the benefits of security to A_1 are only $p(1-q)L = 95$ which is less than the cost of the protective measure. So A_1 will **not** want to invest in protection. In other words, either both divisions invest in security or neither of them does so. These are the two Nash equilibria for this problem.

The Multi-Division Case

The results for the two-division case carry over to the most general settings with some increase in complexity. The incentive for any agent to invest in protection depends on how many other agents there are and on whether or not they are investing. Other agents who do not invest reduce the expected benefits from one's own protective actions and hence reduce a division's incentive to invest. In this section we review briefly the main features of the general case, without providing detailed proofs of the results.[17]

Suppose there are now n identical divisions. Each has a probability p of a catastrophic accident that can cause bankruptcy to the division and $q \leq p$ that the accident bankrupts BeSafe if it does **not** invest in security systems. This probability is zero if it invests. If a division does not invest in protection then it can contaminate all the other divisions. Let $X(n,j)$ be the total expected negative externalities due to contamination from others imposed on a division which has invested in security when j of the other $n-1$ divisions have also taken this step.

If none of the other divisions are protected, then the condition for any division to invest in protection is given by the following condition: $c<p[L-X(n,0)]$. Let c^* be the value of c where the division is indifferent between investing and not investing in protection when j of the other divisions have invested in security. In other words $c^*= p[L-X(n,j)]$. If there are no negative externalities because all the divisions have invested in security, then $c^*= pL$ which is the same as if the division were operating in isolation. As more divisions do not invest in protection then c^* decreases, so that the division is less likely to take security measures if it is maximizing the expected returns of its employees.

[17] The proofs can be found in Heal and Kunreuther (2002).

Now suppose that the number of divisions (*n*) in the organization gets large. When none of the other *n-1* divisions invest in protection then the negative externalities to a division that has installed protection has been shown by Kunreuther and Heal (2003) can be shown to approach $X(n,0) = L$. This implies that the expected loss to any division approaches L as a result of contamination from all the other unprotected divisions. In this situation $c^*=0$ and there is no cost incentive for any division to invest in protecting itself against a catastrophic accident.

Here is the intuition for this somewhat surprising result. If one division is unprotected, then if it incurs a catastrophic accident, there is a probability q that the firm will go bankrupt. One weak link in the organization compromises all the other divisions. In other words, one unprotected division endangers all of the other divisions in the firm even if they have all invested in security. The more divisions that have not invested in protection, the greater the chances that the employees of any division will be looking for another job even if its own plants are secure from a catastrophic accident. As more divisions decide not to invest in security (i.e. take action **N)**, the probability of a catastrophic accident gets very large and there is no economic incentive for your division to undertake protection. In the limit this probability approaches 1 and $c^*=0$.

Tipping and Cascading

In the above examples, all of the divisions in the firm were identical so that there was symmetry with respect to the degree of contamination from one to the others. In reality, there will be some divisions in an organization that may produce much greater negative externalities on the other divisions by their actions than other units. For example, a large division that went bankrupt would be much more likely to cause other divisions to follow suit than a smaller unit in the organization. The large division could suffer a catastrophic loss from an accident that would have much more serious repercussions than if the accident occurred at a smaller plant. By providing incentives for the large division to invest in protection one may convince others in the organization to do the same.

Heal and Kunreuther (2005a) construct an illustrative example involve three agents, two of which are identical and a third one that has a much higher likelihood of creating negative externalities to the other two. In the context of the above example, if Division 3 had a catastrophic accident it would be much more likely to bankrupt the other two divisions than if either Division 1 or 2 had suffered a severe loss. Yet if Division 3 changes from not investing to investing – perhaps as a result of a financial incentive or pressures from the CEO, then both other divisions will voluntarily change as well and there is a new equilibrium at which all are investing in security. The expected profits are increased for each of the divisions so that industry profits are also higher. In fact, Divisions 1 and 2 could afford to subsidize Division 3 so that it would want to invest in security and improve their own expected profits after incurring this additional cost.

More generally, if there are heterogeneous units in the network and there is a weak link that can cause severe disruptions to others, it may only be necessary to provide economic incentives to this unit in order to improve its profitability as well as all of the others in the system. It is this weak link property that characterizes many practical problems in interdependency.

The above model makes the extreme assumption that if a division has an accident it can cause the entire firm to go bankrupt. In reality, an accident from one division that has negative spillover effects on the economic well-being of other divisions will discourage them from undertaking protective actions. More generally, weak links in the system can have major impacts on others even if they have taken action themselves.

It should now be clearer why issues of interdependency require one to have an understanding of the nature of the risks facing each of the chemical plants in an organization as well as the ways that managers deal with the risks. As BeSafe collects more accurate information on the risks of chemical accidents, it can develop more effective strategies for planning at its different divisions.

In their groundbreaking studies on chemical accidents using the RMPInfo database provided by the Environmental Protection Agency (EPA), Kleindorfer and his colleagues provide insight into the nature of the risks associated with accidents from various chemical plants. As detailed in their recent papers in *Risk Analysis* (Kleindorfer et al., 2003; Elliott et al., 2003), they collected data in 1999-2001 on more than 15,000 facilities in the United States that store or use toxic or flammable chemicals believed to be a hazard to the environment or human health. Their database contains the following information:

- The characteristics of the facility itself, including facility location, size and the type of hazard present (as characterized by the chemicals and process used, the training and management systems in place, and other facility-specific characteristics)

- The nature of regulations in force that are applicable to this facility and the nature of enforcement activities

- The level of pressure brought on the facility to operate safely and to inform the community of the hazards it faces;

These data are key inputs to assessing the risk of chemical accidents and better understanding their direct impacts in the form of damage to the plant and the surrounding area as well as the indirect impacts in the form of business interruption losses to the firm as well as other parts of the economy. A detailed analysis of the RMPinfo database as proposed by Elliott et al., (2003) could be extremely helpful in this regard.

Even after such a risk assessment is undertaken, BeSafe and other firms will still have to contend with issues of risk perception at all levels of the organization. Consider the example in Table 1 where each division knew that the probability of a chemical accident was p. It is here where the issue of risk perception comes into play and a manager has to make a decision as to whether or not to invest in protection. Individuals and firms often utilize simplified decision rules when determining what courses of actions to follow. One rule that is often followed by managers is a threshold model of choice, whereby protective action is taken only if the chance of some event occurring is above a critical probability level. If the perceived probability is below this level than it is assumed that this event is not worth worrying about and hence it is not worthwhile to incur the costs of protection.[18]

Empirical research has provided evidence that decision-makers use *threshold-like models* in making their decisions. In a laboratory experiment on purchasing insurance, many individuals bid zero for coverage, apparently viewing the probability of a loss as sufficiently small that they were not interested in protecting themselves against it. (McClelland, Schulze and Coursey, 1993) Similarly, many homeowners residing in communities that are potential sites for nuclear waste facilities have a tendency to dismiss the risk as negligible. (Oberholzer-Gee, 1998). In interview with managers at the Chemco Company after the Bhopal disaster, Bowman and Kunreuther (1988) discovered that the firm had an informal rule that it would only be concerned with the dangers of specific chemicals if the probability of an accident was above a certain threshold.

Suppose a manager utilizes the following rule for its division: Invest in protection only if the probability (p) of having a catastrophic accident that causes its division to go bankrupt is greater than some prespecified value. One could estimate what this threshold probability would have to be so that the division's behavior would be consistent with maximizing the expected returns of its employees.

Consider Division 1 of the firm. Let n^* be the number of other divisions who do **not** invest in protection against catastrophic accidents. One can determine how this threshold probability would be affected as n^* changes. Define p_{n^*} to be the probability that Division 1 would be indifferent between investing and not investing in protection if it were maximizing the expected returns of its employees.

This rule implies that as more divisions do **not** invest in protection, the reduction in the probability of BeSafe going bankrupt by Division 1 investing in protection gets smaller. This means that the value of p_{n^*} that makes it worthwhile for Division 1 to invest in protection must increase for any given c/L ratio. Hence, the larger the value of n^*, the less likely any division will invest in protection for any given value of p. The division thus uses the same line of reasoning in specifying a threshold probability as when it maximizes the expected returns of its employees: investing in protection does less good for the division, the more divisions that have their plants unprotected against a catastrophic accident that could bankrupt the firm.

[18] See Kleindorfer and Klein (2003) for an example of how insurers utilize this type of a threshold model to determine what risks they will want to consider in determining the size and composition of their portfolio.

4. Developing Risk Management Strategies[19]

If divisions within firms are reluctant to adopt protective measures to reduce the chances of catastrophic accidents there may be a need for some type of coordinating mechanism within the organization through internal rules and procedures complemented by private-public partnerships. Both of these options are discussed in this section.

Internal Organizational Rules

In a large decentralized firm with many divisions there is likely to be a need for some type of coordinating mechanism from top management if each division's objective is to maximize the expected returns of its own employees. A key question in this regard is how companies who advertise, "Safety is our most important product" actually operationalize this slogan. Larger firms in the chemical industry have formed functional units that play this role across the organization. For example, DuPont has a process safety management group that is responsible for making sure that all the different divisions in the firm follow appropriate procedures.

In the context of the above example, suppose that BeSafe set up such a cross-cutting unit and instituted a specific rule that would require the division to invest in protective measures when the expected benefits to the firm exceeded the costs of the measure. One way to determine what type of rule to enforce is to consider catastrophic accidents that caused losses so large that it would threaten the solvency of the firm but where the division itself would not want to incur the costs of investing in the measure.

Suppose each of the n divisions had assets of A, and that a loss L to any division was so large that $L>nA$ then the firm would become insolvent should such an event occur. If the cost of protection was c and the probability of an accident was p¸ then the process safety division would always want to require such investment if $c < pL$ to maximize expected profits for the firm.[20] This simple rule effectively eliminates the negative externalities that might otherwise discourage a division from investing in protection. A less intrusive way of encouraging investment in protective measures is to provide subsidies to divisions when $c < pL$ and the division would otherwise not invest in safety.

Role of the Public Sector

A reason for involving the public sector is that some of the consequences of a chemical accident will affect nearby residents but the industrial facility will not be held fully liable for these impacts. For example, suppose there are decreases in property values to homes in the surrounding area or there are disruptions in community life because of an accident. The firm causing the accident will not be legally responsible for these losses.

[19] This section is based on material in Kunreuther and Heal (2005b).
[20] Similarly, if $c>pL$, then the process safety unit would indicate that each division should **not** invest in protection

One way for the government to enforce its regulations is to turn to the private sector for assistance. More specifically, third party inspections coupled with insurance protection can encourage divisions in firms to reduce their risks from accidents and disasters. Such a management-based regulatory strategy shifts the locus of decision-making shifted from the regulator to firms who are now required to do their own planning as to how they will meet a set of standards or regulations. (Coglianese and Lazer, 2001)

The passage of Section 112(r) of the Clean Air Act Amendments (CAAA) of 1990 offers an opportunity to implement such a program. This legislation required facilities to perform a hazard assessment, estimate consequences from accidents and submit a summary report to the U.S. Environmental Protection Agency (EPA) called the Risk Management Plan (RMP) (Belke, 2001). The challenge currently facing the EPA is how to encourage compliance of these regulations so that firms will be operating in a safer manner than they otherwise would be.

There is some urgency for a type of decentralized procedure with appropriate incentives due to the EPA's limited personnel and funds for providing technical guidance and auditing regulated facilities. Chemical firms, particularly smaller ones, have little financial incentive to follow centralized regulatory procedures if they estimate that the likelihood they will be inspected by a regulatory agency are very small and/or they know the fine should they be caught will be low. In such cases they may be willing to take their chances and incur the fine should they violate the existing rule or regulation and be caught. The analogy with the decision on whether to put money in a parking meter seems appropriate. If you know that the chances of a meter being checked are very small and the fine is relatively small then you might think twice before parting with your quarters.

The combination of third party inspections in conjunction with private insurance is a powerful combination of two market mechanisms that can convince many firms of the advantages of implementing RMPs to make their plants safer and encourage the remaining ones to comply with the regulation to avoid being caught and fined. The intuition behind using third parties and insurance to support regulations can be stated rather simply when the regulatory agency has limited personnel to enforce its own rules: low-risk divisions, who the EPA has no need to audit, cannot credibly distinguish themselves from the high-risk ones without some type of inspection.

By delegating part of the inspection process to the private sector through insurance companies and third parties, the EPA provides a channel though which the low-risk divisions in firms can speak for themselves. If a division chooses not to be inspected by third parties, it is more likely to be a high-risk rather than a low-risk one. If it does get inspected and shows that it is protecting itself and the rest of the organization against catastrophic accidents, it will pay a lower premium than a high-risk division which is not undertaking these actions. In this way the proposed mechanism not only substantially reduces the number of inspections the EPA has to undertake, but it also makes their audits more efficient.

Kunreuther, McNulty and Kang (2002) show more formally how such a program could be implemented in practice. They provide supporting evidence from pilot studies undertaken in Delaware and Pennsylvania whereby the Department of Environmental Protection in these two states worked closely with the insurance industry and chemical plants in testing out the proposed program. Similar studies for small firms were undertaken by McNulty et al., (1999).

The process safety management unit of a firm should support this program for two reasons. It gives them a rationale to have the firm hire third party inspectors to make sure their divisions are operating in a safe manner. It also increases the firm's expected profits by reducing the negative externalities that divisions create due to their fear of being contaminated by others.

5. Future Research [21]

The problem of assessing and managing risks associated with chemical accidents highlights the importance of undertaking research on both the descriptive and prescriptive aspects of decision making for low-probability, high-consequence events. Here are some of the areas where there is a need for both basic and applied research.

Risk Assessment

There is a need to collect better data on the risk in order to estimate the chances and consequences of a catastrophic accident. Currently, the Wharton Risk Management and Decision Processes Center is engaging in two types of data collection in this regard with the U.S Environmental Protection Agency. The first of these uses accident history data from the U.S. chemical industry that has already been noted. The second is concerned with the performance of management systems designed to improve the environmental, health and safety (EH&S) performance of companies.

Accident history data can be linked to financial information so one can analyze the association, if any, between the financial characteristics of the parent company of a facility and the frequency or severity of accidents. Similarly, the property damage estimates, and associated indirect costs from these, can be used to assess the consequences of environmental health and safety incidents on overall company performance and provide valuable insights for insurance underwriting for such accidents. Finally, the same data can be used to assess worst-case consequences from such incidents, including those that might arise from site security risks associated with terrorism.

The second data collection project is a study of "near misses" in organizations and the systems that have been put into place to report and analyze these data. (Phimister et al., 2003) Near misses are defined as incidents that, under different circumstances, could have resulted in major accidents. Linking these data on accident precursors to the Accident History database in RMP*Info, may enable one to identify categories of

[21] This section draws heavily on Cohen and Kunreuther (in press).

precursors that give early warnings of the potential for major accidents. Audit tools and other aspects of near-miss management can than focus not just on emergency response but on the range of prevention and mitigation activities before the fact that can help avert major disasters. Even with these data, there will still be considerable uncertainty regarding the estimates of risks associated with these low probability events (National Academy of Engineering 2004).

Risk Perception Issues

There is a need for more detailed information on how the risk interdependencies affect firms' decision processes. Interviews with managers from chemical firms and those in other industries indicate that this feature is normally not taken into account in developing risk management strategies. The recent interest by firms on enterprise risk management suggests that this would be an area that should be incorporated in future analyses.

The models discussed above all assumed that individuals made their decisions by comparing expected benefits with and without protection to the costs of investing in security. There is a growing literature in behavioral economics that suggests that individuals and firms make choices in ways that differ from such a rational model of choice (Kahneman and Tversky 2000). For example, there is evidence that people are myopic and do not appropriately take into account the long-term benefits of investing in protective measures, preferring instead to have a return on their investments over a relatively short time period. It would be useful to understand what factors motivate managers' behavior and to consider strategies for making the investment more worthwhile. Some type of accounting arrangement by the firm to convert the upfront payment into a loan arrangement may enable them to justify the upfront costs while relieving the division of budget constraints that may deter them from making the investment.

It would also be useful to know how managers process information on risk when there is considerable uncertainty on the likelihood and consequences of an accident. We know that individuals have a difficult time dealing with ambiguous risks, particularly those of the low probability variety (Slovic 2000). One telling example is the way the chemical industry behaved prior and after the Bhopal disaster. Prior to the accident, there was a tendency to treat an accident such as the one that occurred in the Union Carbide plant in India as one that will not happen to "our firm." Following the disaster, all chemical companies undertook a detailed study of chemicals with catastrophic risk potential and took special measures to deal with them. (Bowman and Kunreuther 1988).

Risk Management Issues

With respect to managing the risk, there is a need to gain more insight into what impact third-party inspection will have on insurers' willingness to reduce premiums for firms that invest in protective measures. If premiums are regulated to the point that insurers are not permitted to set rates that reflect risks, they are less likely to reward firms for undertaking mitigation measures. However, insurers can still play an important role by certifying that the chemical firm has a risk management plan and is operating in a safe manner. Regulatory agencies and public interest groups will trust the results of the inspection, knowing that insurers are concerned with their own bottom line and would have no incentive to classify a firm as *not risky,* if in fact it posed a high risk. There is, of course, a set of open questions that need to be explored regarding the accuracy of third parties in diagnosing risks that a company faces.

Another area that needs to be examined more carefully is the role that certifications, such as ISO14000, can play in encouraging firms and divisions to operate more safely. In a recent analysis of ISO data and firm performance, Kang (2005) has shown that facilities that have had serious environmental problems are more likely to arrange to be ISO14000 certified than the lower risk facilities and that their performance has improved over the other facilities in the industry after they were certified. There is a tendency for many facilities in a firm to undertake ISO14000 certification procedures at approximately the same time, suggesting that organizations are using this standard as a way of forcing many of their facilities to undergo an inspection that they might otherwise not consider.

Finally, there may be an important role that trade associations can play in providing guidelines for firms to follow with respect to their operations. The Applied Chemistry Council (ACC) [formerly the Chemical Manufacturers Association (CMA)] has undertaken this role through its Responsible Care initiative. Since 1988, members of the ACC have significantly improved their environmental, health, safety and, in recent years, security performance through the Responsible Care initiative. Participation in Responsible Care is mandatory for ACC member companies, all of which have made CEO-level commitments to uphold requirements that includes a management system to drive environmental, health, safety and security performance, sharing progress and activities with the public and having mandatory certification by independent, accredited auditing firms.[22]

6. Issues and Questions for Discussion

The paper conclude with a set of issues and questions to consider with respect to future theoretical and empirical research with respect to IDS problems such as those illustrated in scenarios in Section 2:

Nature of the risk

- What are the similarities and differences between IDS problems? What are their implications for theoretical, empirical, and policy analyses?

- How do IDS examples differ from other problems where there are negative externalities?

---

[22] For more information on the Responsible Care program of ACC see http://www.responsiblecare-us.com/

- How does one incorporate endogenous probabilities into theoretical and empirical analyses? If some agents are known to be more security-conscious, are they less likely to be targets?

- What impact does heterogeneity of agents have on the characterization of the problem and the nature of the solutions? For example, if there are some firms that are much bigger than others, how will this impact on the problem structure and proposed risk management strategies?

Behavioral issues

- Which equilibrium can we expect (if any, and under which circumstances, and for which reasons) in an IDS problem with multiple equilibria?

- How do individuals' biases (such as misperceptions of the risk, the types of heuristics that they utilize in evaluating information, and framing and context) influence their choices in IDS models?

- How do context effects (e.g., terrorism vs. non-intentional hazards) affect individuals' choices?

- How do framing effects (e.g., providing information on the risk in different formats) influence whether an individual will invest in protection?

- How do different forms of communication with others (e.g., computer-mediated interactions, face-to-face discussions) influence individuals' decisions?

Multi-period and dynamic issues

- How does the interaction of network structure and standard learning dynamics influence the outcomes that might actually be selected by a population?

- How do individuals learn when IDS decisions are made over time (in dynamic or repeated games)?

- How does the process of coordination get started?

- What impact do different types of feedback have on the decisions that individuals make over time?

- When do these models lead to tipping and cascading behavior?

Policy instruments and risk-management strategies

- What types of coordinating mechanisms can be designed for improving individual and social welfare? Some of these coordinating mechanisms can be voluntary (e.g., trade associations and compacts) while others can be mandatory (e.g., regulatory actions, well-specified standards).

- Are there ways to create social norms so that individuals want to invest in protection because it is the right thing to do? For example, as smoking became less widely accepted, this increased the incentives for the remaining smokers to quit, and/or modify their behaviors, for example, not smoking indoors.

- Can one use economic incentives such as taxes or fines, perhaps on a few key parties so that they will want to protect themselves and induce others to follow suit? Can the tax or fine be shared with those parties who have already invested in protection? What role can third party inspections and other mechanisms play in the enforcement process?

- How can insurance be used to encourage investment in protective measures when firms are faced with the risk of contamination from other firms?

- What are the appropriate roles of public-sector interventions such as regulations and standards (e.g., building codes) for inducing individuals to invest in protective measures?

REFERENCES

Auerswald, P., L. Branscomb, T. LaPorte, and E. Michel-Kerjan, (2006) *Seeds of Disaster, Roots of Response: How Private Action Can Reduce Public Vulnerability*. New York: Cambridge University Press.

Ayres, I. and S. Levitt (1998), "Measuring the Positive Externalities from Unobservable Victim Precaution: An Empirical Analysis of LoJack" *Quarterly Journal of Economics* 113:43-77.

Barrett, S. *Environment and Statecraft: the Strategy of Environmental Treaty-Making* Oxford University Press, (2003)

Belke, J. (2001), "The Case For Voluntary Third Party Risk Management Program Audits." Paper presented at the 2001 Process Plant Safety Symposium of the American Institute of Chemical Engineers, April 23.

Bowman, E.H. and H. Kunreuther, (1988), "Post Bhopal Behavior of a Chemical Company," *Journal of Management Studies*, 25, 387-402.

Coglianese, C. and D. Lazer, (2001), "Management-Based Regulation: Using Private Sector Management to Achieve Public Goals" *Regulatory Policy Program Working Paper,* Cambridge, MA: John F. Kennedy School of Government, Harvard University

Cohen, M. and H. Kunreuther (in press) "Operations Risk Management: Overview of Paul Kleindorfer's Contribution," *POMS*

Elliott, M.R., P.R. Kleindorfer, and R.A. Lowe, (2003), "The Role of Hazardousness and Regulatory Practice in the Accidental Release of Chemicals at US Industrial Facilities," *Risk Analysis*, 23, (5), 883-896.

Feinstein, J. (2006) "Managing Reliability in Electric Power Companies" in *Seeds of Disaster, Roots of Response: How Private Action Can Reduce Public Vulnerability*. P. Auerswald, L. Branscomb, T. LaPorte, and E. Michel-Kerjan, (eds.) New York: Cambridge University Press.

Greenwald, B. and J. Stiglitz, (1990), "Asymmetric Information and the New Theory of the Firm: Financial Constraints and Risk Behavior," *American Economic Review: Papers and Proceedings* 80:160-165.

Heal, G. and H. Kunreuther, (2002) "The Vaccination Game," Columbia Business School and The Wharton School, July 2002, revised January 2005

Heal, G. and H. Kunreuther, (2005a), "IDS Models of Airline Security" *Journal of Conflict Resolution* 41:201-17.

Heal, G. and H. Kunreuther, (2005b) "You Can Only Die Once" in H.W. Richardson, P. Gordon and J.E. Moore II, (eds.), *The Economic Impacts of Terrorist Attacks.* Cheltenham, UK: Edward Elgar.

Heal, G., M. Kearns, P. Kleindorfer and H. Kunreuther (2006), "Interdependent Security in Interconnected Networks," in *Seeds of Disaster, Roots of Response: How Private Action Can Reduce Public Vulnerability*, P. Auerswald, L. Branscomb, T. LaPorte, and E. Michel-Kerjan, (eds.) New York: Cambridge University Press.

Kahneman, D. and A. Tversky, (2000) *Choices, Values and Frames* New York: Cambridge University Press.

Kang, Y. (2005) "The Impacts of Third Party Inspections on Industrial Safety and Environmental Performance" Philadelphia: Wharton School, University of Pennsylvania (mimeo)

Kearns, M (2005). "Economics, Computer Science, and Policy." *Issues in Science and Technology*, Winter 37-47.

Kleindorfer, P.R., J. Belke, M.R. Elliott, K. Lee, R.A. Lowe, and H. Feldman, (2003) "Accident epidemiology and the US chemical industry: accident history and worst-case data from RMP," *Risk Analysis*, 23(5): 865-881.

Kleindorfer, P.R., M.R. Elliott, Y. Wang and R.A. Lowe, "Drivers of accident preparedness and safety: evidence from the RMP rule," *Journal of Hazardous Materials*, (2004), Vol. 115 (August), 9-16.

Kleindorfer, P. R. and R. W. Klein, "Regulation and Markets for Catastrophe Insurance," in *Advances in Economic Design*, Murat R. Sertel and Semih Koray (eds.), Springer Verlag, Heidelberg, 2003

Kunreuther, H and G. Heal (2003),"Interdependent Security," *Journal of Risk and Uncertainty*, (2003) Springer vol. 26(2-3) pp. 231-49.

Kunreuther, H, and G. Heal, (2005). "Interdependencies in Organizations," in B. Hutter and M. Powers (eds.) *Organizational Encounters with Risk*, Cambridge: Cambridge University Press.

Kunreuther, H., P. McNulty, and Y. Kang, (2002), "Improving Environmental Safety Through Third Party Inspection," *Risk Analysis* 22:309-18

McClelland, G., W. Schulze, and D. Coursey, (1993), "Insurance for Low-Probability Hazards: A Bimodal Response to Unlikely Events," *Journal of Risk and Uncertainty* 7:95-116.

McNulty, P., R.A. Barrish, R.C. Antoff, and L.C. Schaller, (1999), “Evaluating the use of third parties to measure process safety management in small firms.” 1999 Annual Symposium, Mary Kay O’Connor Process Safety Center, Texas A&M University, October 26.

National Academy of Engineering (2004) *Accident Precursor Analysis and Management,* Washington DC, The National Acadamies Press.

Oberholzer-Gee, Felix. (1998). “Learning to Bear the Unbearable: Towards and Explanation of Risk Ignorance,” Mimeo, Wharton School, University of Pennsylvania.

Phimister et al., 2003) “Near-Miss Incident Management in the Chemical Process Industry,” *Risk Analysis*, Vol 23, No. 3. pp 445-459, 2003

Slovic, Paul (2000). The *Perception of Risk.* London, UK: Earthscan.

Spyratos, V., P. Bourgeron, and M. Ghil, (2007) “Development at the wildland–urban interface and the mitigation of forest-fire risk,” The National Academy of Sciences of the USA