

Taking risk away from risk taking: decision insurance in organizations

Tad Hogg and Bernardo A. Huberman
HP Labs
Palo Alto, CA 94304

April 13, 2006

Abstract

We present a new mechanism for encouraging risk taking within organizations that relies on the provision of decision insurance to managers. Since insurance increases the likelihood of free riding, we also introduce a technique that mitigates this moral hazard by automatically identifying the social network around the manager and using it as a monitoring group.

We show that three possible regimes exist. In the first one, managers contribute to production but avoid risky projects. In the second, managers take on risky projects without free riding. In the third, they free ride. We establish the conditions for the appearance of each of these regimes and show how to adjust the mechanism parameters so as to get the highest expected payoff for the firm in spite of its risk-adverse managers.

1 Introduction

While most organizations consider taking risks as an essential part of their success, few have actual policies to encourage risk taking on the part of their managers and employees. Among the reasons, there are subtle assumptions about risk taking that give it a negative connotation within an organizational setting [15]. As shown by March and Shapira [16], the discussion around the value of risk in most organizations is two sided. While it is often stated that risk taking is essential to innovation and progress, it is also sharply distinguished from gambling or playing the odds, which has a negative connotation. Managers emphasize this difference because their experience teaches them that the appropriate choices are those involving undesirable outcomes that can be avoided, whereas by gambling they actually mean risk taking that turns out badly. Thus, risky choices that fail are seen as mistakes that could be avoided, a perception that discourages managers from taking those risks. This results in organizations that reward outcomes, not decisions.

In addition to institutional barriers, there are problems that stem from the way individuals actually evaluate and act when confronted with risk [11]. It has been known for a long time that people's perception of what constitutes a random event does not conform to its statistical definition [29]. Equally important, the evaluation of probabilistic events suffer from framing issues that are by now well documented. Early on, Kahnemann and Tversky [12] established that when individuals deal with risky alternatives whose possible outcomes are generally good they tend to be risk averse, but when dealing with risky alternatives whose possible outcomes are poor they become risk seeking. Moreover, a number of empirical studies have shown that managers are quite insensitive to probability estimates of possible outcomes and do not understand or trust probabilistic measures [26, 4, 27]. And when they can estimate probabilities, they are often reluctant to change their behavior because of their status quo biases [22, 2].

A large empirical study of decision making by managers in different organizations [23] showed that it is often the magnitude of the value of the outcome that determines what managers do when making decisions, rather than its weighting by the probability of its likelihood. In addition, other authors [1, 18] showed that managers make decisions based not on the unpredictability of outcomes but on their costs, which renders risk equivalent to hazard—the value of an extreme outcome rather than its variability.

Furthermore, it is often the case that if a manager is above a performance target

her primary focus is on avoiding actions that might place her below it, thus making the opportunities for gain less salient. And for those below target, the desire to reach it focuses attention in a way that leads to excessive risk taking [14, 13, 23].

Given all these considerations and the vast literature that validates them, it is desirable to establish mechanisms and policies that encourage risk taking on the part of managers that—for the multifaceted reasons outlined above—are risk averse.

A familiar incentive consists in awarding stock options to executives to align their interest with that of the organizations they lead. Stock options can be effective [21], but can also suffer from deleterious outcomes [7]. For example, if priced too high and over too short a term, they can encourage the executive to make risk seeking decisions that can be dangerous for the company, in the hope of making the option valuable to her. On the other hand, if priced too low, options tend to make the executive risk averse, as he will want to keep the advantage conferred on the day of their issue.

Other kinds of compensation schemes include pooling and tournaments. Such compensation schemes can be useful in small teams as they both avoid the free rider problem and provide stronger incentives through a visible connection between effort and results [17]. The work we present here can be seen as similar in spirit to this team based profit sharing approach, with the added advantage of providing an automatic mechanism to identify suitable teams through communities of practice discovered through online network communication patterns [28].

Another possible way to encourage risk taking would be to resort to new kinds of insurance. Recently, several economists have suggested extending insurance to a variety of other risky endeavors, such as consumption-indexed public pension funds [19], options on business cycle variables such as consumer confidence indices [10], the choice of a professional career and volatile property values [24, 25]. While provocative, these proposals suffer from the problem that their implementation requires the creation of groups willing to assume these risks and to gain enough experience for actuarial tables to become useful tools in assessing risk and setting premiums. Equally difficult, the implementation of any new insurance mechanism has to contend with moral hazard, i.e the tendency of individuals to change their behavior in such a way that the insured outcome becomes more probable. Another issue with insurance is adverse selection, in which only the poor performers choose to seek insurance.

As we show in this paper, a properly designed decision insurance mechanism

can be readily implemented within organizations so as to increase the likelihood that managers will behave towards the best interest of the enterprise even when they are risk averse. Because in general there is no one-to-one mapping between a decision and an outcome compensation schemes solely based on outcomes have difficulty discriminating among the decision choices that managers had that led to those outcomes. This leads to risk averse managers emphasizing decisions that lead to a subset of safe outcomes.

Equally important, we show how one can use the social network of the manager to help reduce the moral hazard inherent in this insurance. We do so by automatically identifying the community of practice around a given manager from patterns in email exchanges in the organization. From this community of practice one can set up a monitoring group whose ability to detect free riding more than compensates the possibility of false positives, ie. construing as free riding what is not on the part of the manager. Likewise, this implementation within the enterprise offers the possibility of addressing adverse selection by making participation mandatory for all managers.

A quantitative analysis of all the factors that enter into this decision insurance leads to the conclusion that three possible behavioral regimes exist. In the first one, managers contribute to production but avoid risky projects. In the second regime, managers adopt a risk neutral stance when taking decision insurance. In the third one, they free ride by not giving the full attention to the decision problem at hand while relying on their share of revenue from the efforts of others to compensate for the resulting low performance of their own projects.

These three regimes are illustrated by considering a model of risk aversion based on expected utility. While this model does not capture the full range of human behavior under risk, it nevertheless gives useful quantitative insights on how decision insurance can operate in realistic organizational settings. We then compute the conditions for the appearance of each of these regimes and show how one can adjust the payoffs so as to be in the desirable state, i.e., giving the highest expected payoff for the organization in spite of having risk-averse managers.

This paper is structured as follows. We first discuss decision insurance and present several perspectives on it that allow for a quantification of the premiums and payoffs that are at stake. We also show how insurance makes risk averse individuals behave in an almost risk-neutral fashion.

We then proceed to an analysis of moral hazard in the context of organizational

decision making and show how a monitoring group made up of members of the manager's community of practice can be structured so as to increase the likelihood of detection of free riding while minimizing false positives. The next section describes the algorithm used to discover a community of practice around a decision maker from patterns of email communication.

Finally the equilibrium outcomes are computed and the dynamics that such a mechanism can induce in the organization is discussed. We conclude by presenting the concrete steps that are necessary to implement decision insurance within an organization and the problems that are likely to be encountered when such a mechanism is first deployed.

2 Insurance

Insurance is a commonly used mechanism to hedge against risk, whereby a group or company assumes the risk in return for premium payment. This works well when the number of insured people is large and their risks are independent of each other, as the many examples of shipping, casualty and home insurance have shown over the years.

From an organizational perspective, managers should choose those projects that give the highest expected value to the enterprise. When risky projects do have a higher expected value, risk averse managers will be reluctant to choose them in spite of their having a higher expected return. The instantiation of decision insurance mechanism will allow them to transfer their risk to the organization.

From the individual decision maker's perspective, decision insurance amounts to her paying a premium in exchange for a payoff that is larger than the premium when the project he decides to undertake fails. When the project succeeds the decision maker receives instead a bonus also larger than the premium and proportional to the value to the company of the successful project.

notation	meaning
V_{risky}	payoff when risky project succeeds
p	probability risky project succeeds
V_{safe}	payoff of safe project
q	probability group member gives accurate evaluation
Q	probability group reaches threshold for reporting free riding
q_{false}	probability group member incorrectly detects free riding
Q_{false}	probability group incorrectly reaches threshold

Table 1: Symbols used to express expected payoffs. The first three rows refer to the decision payoffs, while the last four relate to the moral hazard.

Consider first an organization whose decision makers have a choice between risky and safe projects with different returns to the enterprise. As an incentive for success, managers receive some portion of those returns. This results in payoffs to decision makers based on the outcomes of their decisions. Suppose that a successful risky project has an expected payoff that is r times the payoff of a safe project. This means that the organization can afford to pay on average up to r times whatever it takes for people to do the safe project and still be better off.

We look first at a safe project which always produces a certain return to the company. Of this, V_{safe} is the amount the organization sets aside to provide extra compensation (bonuses, stock options, etc.) to its managers. Similarly V_{risky} is the amount made available for extra compensation if the risky project are successful. Assume further that the risky project succeeds with probability p . Thus the expected amount is $x = pV_{\text{risky}}$. These parameters for the decision payoffs are summarized in Table 1.

This extra compensation can be paid to managers in two possible ways. First, each manager gets extra pay-for-performance compensation resulting from their own decisions. This means that they get either V_{safe} or either 0 or V_{risky} , depending on which project they choose. Second, in the case where they used decision insurance, they receive the payments from the compensation pool according to our mechanism. This consists of a bonus that is a share of the compensation pool.

When insurance is in place and no employee free rides, each manager is better off by choosing the project with the higher expected return. This is because risk aversion on the part of manager is less pronounced when small rewards are at stake. The outcome is that everyone then takes on the risky project. Therefore profits are r times larger than in the case of no insurance and if the payment is in the form of profit sharing, $1/n$ of the employees get r times what they would have otherwise gotten. Notice that when the project is successful those that took this form of insurance end up being compensated less than they would have if they chose not to take insurance while taking the risky project.

In order to relate these two views of the mechanism let the premium that a manager pays to the organization be m , in exchange for a payment $V_{\text{insurance}}$ if the project fails. Thus, with probability p the manager receives $(V_{\text{others}} + V_{\text{risky}})/n$ and with probability $1 - p$ he receives V_{others}/n . Here V_{others} is the total payoff from other decision makers' projects. In traditional terms these correspond to receiving the benefits of the project less the premium plus any payoff from the insurance. Thus, when the project is successful the manager receives $V_{\text{risky}} - m$ and when it is not successful $V_{\text{insurance}} - m$. Equating these two sets of expressions for the outcomes gives

$$V_{\text{insurance}} = V_{\text{risky}} \left(1 - \frac{1}{n}\right) \quad (1)$$

$$m = V_{\text{insurance}} - \frac{1}{n}V_{\text{others}} \quad (2)$$

To connect with traditional fixed value premiums and payoffs when the group is

large, V_{others} will be nearly its expected value when all other decision makers work on projects with the same expected payoff.

Another perspective on this approach arises from notions of utility and risk aversion. A common technique used to characterize risk attitudes is to consider the function U which gives the utility to the individual as a function of the dollars received or lost. In particular, a concave utility function means that the individual prefers a sure payoff to an uncertain one with the same expected value. The concavity of the curve corresponds to the amount of risk aversion and relates to the degree to which the individual values an extra dollar as his income goes up. While widely used in economics, the notion of utility maximization of a concave utility function as an explanation of risk aversion is not without problems [20], as it tends to predict unreasonable behaviors for large value ranges. Nevertheless, for a small range of values, it provides a reasonable model of risk aversion [5]. In the case of insurance within an organization, for managers who make many decisions which are relatively compared to their lifetime earnings, the amounts involved are small, thus making the model reasonable. Moreover, insurance is well known to mitigate risk in practice. Thus it applies to the actual risk behavior of people, regardless of whether this risk is well described by the utility model we use.

The effect of insurance on an individual's risk aversion is to reduce the variation in the range of dollars received for the different outcomes. In this smaller range the concavity of the utility curve is less significant, causing the behavior to be closer to a risk neutral one. We illustrate this behavior schematically in Fig. 1. To do so, we pick a particular form for the utility function, i.e.

$$U(V) = \ln(V + W) \quad (3)$$

where W the initial wealth of the individual and V is the number of dollars received. This function is concave, representing the risk averse nature of the individuals. The preference for a certain outcome compared to an uncertain one with same expected value decreases for larger value, representing the diminishing benefit from an additional dollar the more one has.

As can be seen in the figure, without insurance the safe project, A, has a higher utility even though it has a lower expected value. Thus the risk averse manager will produce lower expected returns to the organization on average. With insurance however, the expected utility from the risky project is thus higher. In this case the insurance is sufficient to shift the decision making to being the same as the one a risk neutral manager would pick.

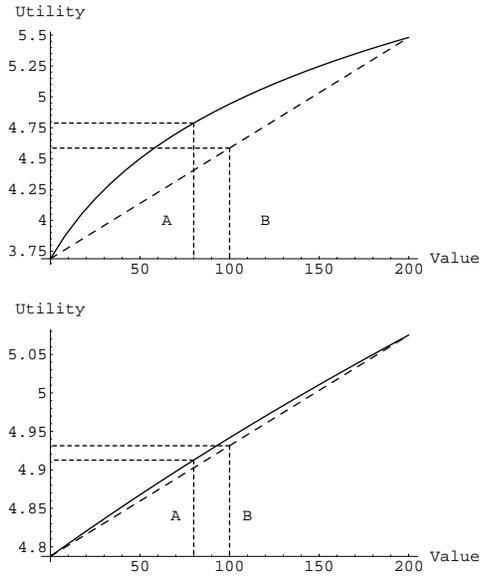


Figure 1: Utility function for a risk averse individual as a function of dollars received for the individual project with and without insurance. We use the utilities given by Eq. (3) with $W = 40$ and $V_{\text{safe}} = 80$, $V_{\text{risky}} = 200$, $p = 0.5$. Both figures compare the utilities of the safe project (A) with the expected utility of the risky project (B). The first shows the case without insurance in which the individuals receive the payoffs of their individual projects. The second figure shows the case with insurance in a group of size 5, where everyone else is choosing the project with the higher expected utility, i.e. the risky one.

So far we have assumed that all the managers's risk attitudes and skills are the same. In real life however, there is variation in these attributes. Decision insurance can adjust for this by having different ratios of pay for performance to profit sharing for different managers, which amounts to having different premiums. This is important to realize, for otherwise adverse selection will take place, i.e. only the least skilled managers will tend to join the enterprise.

2.1 Moral hazard

A significant problem with this kind of insurance, both outside the enterprise [25] and inside it, is moral hazard, i.e., individuals changing their behavior because they no longer bear the full cost of their decisions. In our organizational context, this arises through the effort that managers exert into evaluating the likelihood of given outcomes. With compensation no longer tied to the success of the individual's efforts, there is a temptation to work less and thus suffer only a small decrease ($1/n$) in individual compensation, i.e. free ride on the effort of others in the organization.

Thus the desire to encourage risk-taking can create a moral hazard problem for the organization. This hazard occurs when the reduced connection between individual effort and reward tempts managers to work less diligently. Even if group members don't actually follow such a strategy, the appearance of a conflict of interest caused by insurance could inhibit its widespread adoption, and thus prevent the organization from realizing its potential for improved performance from risk-taking.

This is an example of the more general agency problem in which one group is unable to fully monitor the efforts of another and so needs to align as much as possible the incentives of both groups [3]. Agency theory examines relationships in which one party (the principal) determines the work, which another party (the agent) undertakes. The theory argues that under conditions of incomplete information and uncertainty, which characterize most business settings, two agency problems arise: adverse selection and moral hazard. Adverse selection is the condition under which the principal cannot ascertain if the agent accurately represents his ability to do the work for which he is being paid. And moral hazard is the situation under which the principal cannot be sure if the agent has put forth maximal effort.

The free riding problem is conventionally addressed through monitoring or reputation mechanisms. In an organizational context, managers provide some level of performance monitoring, but may not be in a position to know whether outcome likelihoods used to make decisions are accurately estimated or simply determined arbitrarily in the knowledge that the insurance mechanism will compensate for any failures.

To improve the accuracy of decisions (and by implication encourage a substantial effort on the part of decision makers) we propose supplementing manager

oversight with a monitoring group of coworkers who are likely to be familiar with the decision maker's choices. This group can also provide occasional advice to the decision maker as to the likely outcomes of his decisions ¹.

Specifically each member of the group assesses whether or not the manager is free riding in the context of a decision making event. If at least t out of the N members of the group detect free riding, then upper level management can decide to investigate and whether or not to impose a penalty, C on the decision maker. This penalty could take the form of an increased premium or other sanction. The size of the penalty should balance the need to deter free riding and the possibility of false positives, which tend to make people hesitate to take risky projects.

In designing this monitoring mechanism one has to select suitable threshold values. To see how this assessment would work, consider the case where the insured is free riding. What is the chance that this mechanism will detect it? If each member of the monitoring group can detect this free riding with probability q , and assuming independent evaluations, the probability Q that this threshold is achieved is given by the upper tail of the binomial distribution

$$Q(t, q, N) = \sum_{i=t}^N \binom{N}{i} (1-q)^{N-i} q^i \quad (4)$$

Eq. (4) gives the probability that the aggregated group estimates that the decision maker is free riding, thereby indicating a potential problem. On the other hand, there is the possibility of false positives, i.e. a member of the monitoring group detects free riding when there is none. Suppose this happens with probability q_{false} . Then the probability of a false positive for the group as a whole is Q_{false} given by Eq. (4) with q_{false} instead of q . These parameters for the moral hazard are summarized in Table 1.

Ideally one would like to find a threshold large enough to avoid false positives, while still having a reasonable chance of detecting free riding when it occurs. Whether this is possible depends on the ability of group members to discriminate among these situations².

¹This group is reminiscent of the friendly societies common prior to modern insurance companies, where people of similar social background provided mutual support while being able to monitor any free riding.

²An interesting issue with the choice of a monitoring group is whether or not the ability to discriminate between these cases depends on the distance from a member of the group to the insured

If group members have a large ability to discriminate, then $q_{\text{false}} \ll q \approx 1$. On the other hand, if members lack this ability these values will be nearly the same. In the first case it will be possible to choose the threshold so as to have a high group probability of detecting free riding with very low false positives. Otherwise, it won't be possible. This is illustrated in Fig. 2. As can be seen when individuals can discriminate well between actual free riding and appearances, it is possible to select the threshold so that $Q \approx 1$, while the probability of a false positive is very small.

One may wonder if the existence of a peer monitoring group might act so as to encourage those who are take insurance to behave in a way that they expect the majority of the group to agree with. If so, this could counteract the benefit of the insurance when the majority is more risk averse than the insured. Thus although the monitoring group can reduce free riding, care has to be taken that it does not adversely affect the decision making of the insured.

Choosing a high threshold will ensure that the formal mechanism is rarely exercised while its mere existence and the desire to maintain one's reputation (the insured) with his peer group will act to ensure dedication to delivering on the outcome of the selected project. With these features/choices, monitoring groups can encourage correct behavior without having to act very often.

within the social network. Our assumption is that those closer to the insured have a better ability to detect accurately the intent of the free rider.

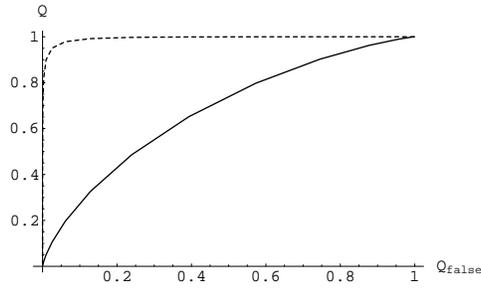


Figure 2: Behavior of Q vs. Q_{false} as the threshold varies from 0 to N . The curves show cases with high (dashed curve) and low (solid curve) discrimination on the part of the group members. When the threshold is low both detection probabilities are small, in the lower left hand corner of the plot. As the threshold increases, both probabilities increase. In this figure $N=30$, $q_{\text{false}} = 0.2$ for both curves, and $q = 0.25$ for the solid curve, and $q = 0.5$ for the dashed curve.

2.2 Finding the community of practice

In order to establish a group of individuals that can monitor and verify that the insured engage in behavior that is beneficial to the enterprise one needs to identify those who have a familiarity both with the manager and also with the nature of the work that he engages in. Since it is usually those belonging to the manager's community of practice that have those properties, an obvious strategy is to determine those communities of practice from the social network inside the organization.

The advent of email as the predominant means of communication in the information society offers a unique opportunity to observe the flow of information along both formal and informal channels. Not surprisingly, email has been established as an indicator of collaboration and knowledge exchange, as well as the speed with which information travels within organizations [8]. This volume of data enables the discovery of shared interests and relationships where none were previously known, as well as ways to automatically identify communities of practice from the social network of an organization [28]. A particularly effective one relies on the examination of email patterns within an enterprise [28, 30] and can clearly and swiftly (without analyzing the contents of the messages) reveal members of a community of practice to which anyone belongs [30].

The automatic discovery of communities from email patterns works by first creating a graph whose nodes are individuals and the edges signify connections

among them. These connections are determined by the frequency with which any pair of individuals exchanges email over a particular period of time, and can be arbitrarily set by the mechanism designer. Once such a graph is constructed for the whole or part of the organization, one can quickly find the community structure within the graph. By community structure within a graph we mean that the graph can be divided into groups so that edges appear within a group much more often than across groups.

Recently, Wu and Huberman[30] introduced a fast method that allows for the discovery of communities within graphs of arbitrary size in times that scale linearly with the size of the graph. This method, which relies on treating the graph as an electric circuit, is based on notions of voltage drops across networks that are both intuitive to grasp and easy to solve regardless of the complexity of the graph involved. Most importantly for our purposes, the method allows for the discovery of a community associated with a given node without having to extract all the communities out of a graph.

To see how the algorithm works, consider the simplest problem, that of dividing a graph into two communities. Consider a graph $G = (V, E)$ and assume for simplicity that we already know that node A and B belong to different communities, which we call G_1 and G_2 . The basic idea is to imagine each edge of this graph to be a resistor with the same resistance, and to connect a battery between A and B so that they have fixed voltages, say 1 and 0. Having made these assumptions, the graph becomes an electric circuit with current flowing through each edge (resistor). By solving Kirchhoff equations one can then obtain the voltage value of each node, which should lie between 0 and 1. The claim of the method is that, from a node's voltage value, one can judge whether it belongs to G_1 or G_2 . Specifically, one can say a node belongs to G_1 if its voltage is greater than a certain threshold, say 0.5, and that it belongs to G_2 if its voltage is less than that threshold.

First let us consider the simplest case that node C has only one neighbor D , so logically C should belong to the same community as D . Because no current can flow through the edge CD , the two endpoints must have the same voltage, thus they belong to the same community.

Next, consider the case where node C connects to two neighbors D and E . Because the edges CD and CE have the same resistance, one must have $V_C = (V_D + V_E)/2$. Hence if D and E belong to the same community, i.e., V_D and V_E both lie above or below the threshold, then V_C lying between V_D and V_E should be above or below the threshold as well, therefore belonging to the same community

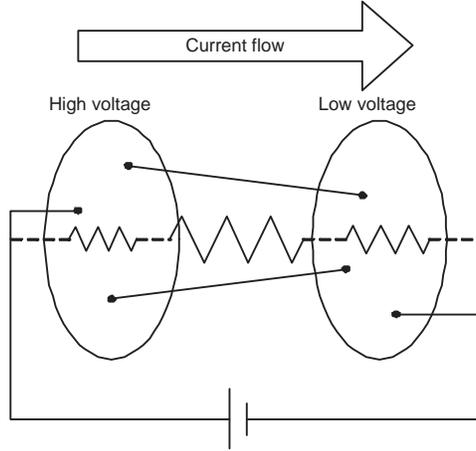


Figure 3: A network with a voltage applied across two nodes. Current flows from left to right, thereby building a voltage difference. Because nodes inside a community are densely connected, their voltages tend to be close to each other. A big voltage gap happens about halfway between the two communities, where the edges are sparse and the local resistance is large.

as D and E , which makes sense. On the other hand, if D and E belong to different communities, then it is comparatively hard to tell which community C belongs to (V_C might be near the threshold), but this is exactly where ambiguity arises - a node has connections with more than one communities.

Lastly we consider the most general case: node C connects to n neighbors D_1, \dots, D_n . The Kirchhoff equations tell us that the total current flowing into C should sum up to zero, i.e.,

$$\sum_{i=1}^n I_i = \sum_{i=1}^n \frac{V_{D_i} - V_C}{R} = 0, \quad (5)$$

where I_i is the current flowing from D_i to C . Thus

$$V_C = \frac{1}{n} \sum_{i=1}^n V_{D_i}. \quad (6)$$

Thus the voltage of a node is the average of its neighbors. If the majority of C 's neighbors belongs to a community which has voltage greater than the threshold,

then V_C tends to exceed the threshold as well, and this method will classify C into that community.³

This methodology, supplemented with repeated random choices of any two nodes in a large graph, is extremely efficient at discovering the community structure of any social network in times that scale linearly with the size of the community. Equally important, it allows to find the community of a given node instead of all communities. To this end, instead of randomly picking two nodes at a time, one fixes the given node/individual as one pole, and chooses the second pole to be another random node that is at least a distance 2 away from the first one.

This concludes the identification of the monitoring community. What now remains is the determination of the particular parameter values that make for a useful decision insurance mechanism, ie. one that both promotes risk taking by managers while avoiding the moral hazard of their free riding on the efforts of others.

³This method can be easily extended to weighted graphs. All we need to do is to set each edge's conductivity proportional to its weight:

$$R_{ij} = w_{ij}^{-1}. \tag{7}$$

3 An example

To understand how this mechanism works, consider its application to managers with the same preferences and having three choices:

- free ride by exerting less effort at assessing likely outcomes with a baseline cost (which we take to be 0) to the manager and a baseline benefit to the group
- choose the safe project with cost c to the manager and payoff value V_{safe} to the group
- choose the risky project, also with cost c to the manager, with variable payoff, which for simplicity we take to be V_{risky} with probability p , and 0 otherwise

In the case of free riding there is the additional possibility of getting caught with probability Q , in which case the manager pays a penalty C . The probability Q depends on the structure of the manager's social network, the choice of threshold, and the ability of the group members to detect significant missed estimates, according to Eq. (4). For the sake of simplicity this example will assume high discrimination ability on the part of the monitors so that the false positive rate is negligible. If this were not the case it amount to an extra contribution to the expected cost incurred by the manager. This extra contribution reduces the difference in payoffs between free riding and not free riding but not the qualitative regimes exhibited in this example.

With these parameters, summarized in Table 1, we can denote the expected values of the three choices in the two situations. Without insurance they are given by

$$E_{\text{free-ride}} = U(0) \quad (8)$$

$$E_{\text{safe}} = U(V_{\text{safe}} - c) \quad (9)$$

$$E_{\text{risky}} = (1 - p)U(-c) + pU(V_{\text{risky}} - c) \quad (10)$$

whereas with insurance they become

$$E_{\text{free-ride}}^{\text{insurance}} = (1 - Q)U\left(\frac{V_{\text{others}}}{n}\right) + QU(-C) \quad (11)$$

$$E_{\text{safe}}^{\text{insurance}} = U\left(\frac{V_{\text{safe}}}{n} - c + \frac{V_{\text{others}}}{n}\right) \quad (12)$$

$$E_{\text{risky}}^{\text{insurance}} = (1-p)U\left(\frac{V_{\text{others}}}{n} - c\right) + pU\left(\frac{V_{\text{others}}}{n} + \frac{V_{\text{risky}}}{n} - c\right) \quad (13)$$

The situation we will consider is when decision insurance is relevant, i.e., the expected value of the risky project exceeds the value of the safe project, i.e. $V_{\text{risky}}p > V_{\text{safe}}$. In this case the maximum expected value for the group occurs when managers choose the risky projects.

How managers behave when confronted with these changes depends on their risk attitudes. Without decision insurance, we assume that people prefer not to free ride, so that

$$E_{\text{safe}} > E_{\text{free-ride}} \quad (14)$$

While the risky project gives the highest expected payoff to the manager (i.e., $V_{\text{risky}}p - c$), which is larger than what he gets from the safe project, (i.e., $V_{\text{safe}} - c$), this is not necessarily true from the point of view of the manager's utilities. Specifically, assume that managers are sufficiently risk adverse so that they prefer the safe project over the risky one, i.e.

$$E_{\text{safe}} > E_{\text{risky}} \quad (15)$$

Since this situation has a higher expected benefit for the organization than when risky projects are picked, we must require that $pV_{\text{risky}} > V_{\text{safe}}$.

With decision insurance in place, the payoff to a manager is $1/n$ of the payoff to the group (profit sharing). Notice that when the expected benefit from either risky and safe projects are both less than the benefit of free riding, managers will be tempted to free ride and not contribute to the profits. This amounts to having

$$U\left(\frac{V_{\text{others}}}{n}\right) > E_{\text{safe}}^{\text{insurance}} \quad (16)$$

$$U\left(\frac{V_{\text{others}}}{n}\right) > E_{\text{risky}}^{\text{insurance}} \quad (17)$$

When the group is large, this latter criterion amounts to $pV_{\text{risky}}/n < c$.

If decision insurance is now combined with anonymous monitoring sufficient to make the cost of free riding higher, the choices left to managers are among the safe project and the risky one. In order for managers to choose the risky project the final relation one needs is

$$E_{\text{risky}}^{\text{insurance}} > E_{\text{safe}}^{\text{insurance}} \quad (18)$$

3.1 Equilibrium outcomes and adaptive behavior

For simplicity, we will first assume that all the organization managers have identical preferences. In this case, if there is a simple equilibrium all will make the same choice, and the fixed point will correspond to the choice with the highest expected utility.

To illustrate the possible behaviors, we use the functional form for the utility given by Eq. (3). This allows us to explicitly illustrate the range of behaviors that can ensue as a function of the parameter values.

The resulting regimes are shown in Fig. 4. As can be seen, for these parameter values one can obtain several equilibria, separated by the curves shown. For lower values of Q and the full range of probabilities for the risky project to be successful, managers will free ride. For larger values of Q there are two possible outcomes. Either everyone chooses the safe project or the risky. When there is no insurance the transition between these occurs at $p = 0.54$, whereas with insurance this transition is lowered to $p = 0.404$, which is just slightly above the risk neutral location of 0.4.

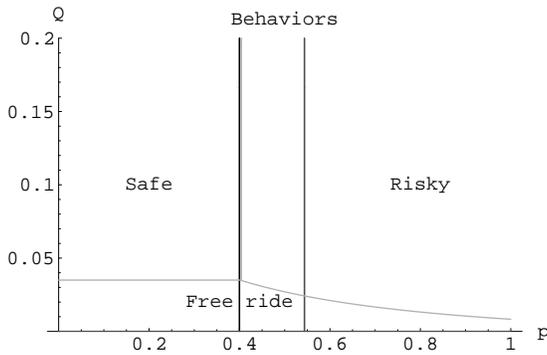


Figure 4: Equilibrium regimes for parameters $V_{\text{safe}} = 80$, $V_{\text{risky}} = 200$, $c = 10$, $C = 50$, $n = 35$, $W = 100$. The curves denote the boundaries between the behaviors with the highest utilities. The vertical line on the right corresponds to no insurance, whereas the left one is the case with insurance, which is indistinguishable from the boundary for perfect risk neutral managers. The horizontal curve denotes the boundary between free riding and no free riding.

While this results illustrate the kind of equilibria that can ensue from the existence of decision insurance it is also important to discuss the dynamics that lead

to these fixed points. It is often the case that as external conditions in the market-place change, those changes are detected by those most directly involved in dealing with customers. These changes tend to alter the perception of the payoffs and the likelihood of the success of particular projects. The question then is: how does decision insurance respond to these changes, or, specifically how do managers respond given the incentives offered by the mechanism? In particular, we need to answer to what extent it can respond correctly without needed further intervention or retuning by management? By this we mean adjusting the choice among risky and safe project while preventing free riding.

Consider a scenario whereby initially all the managers of the organization choose to work on safe projects because their assessment of the success probability of a risky project is low. Over time additional information may change this assessment and therefore the alternative has a greater expected utility. If this probability continues to increase, a typical manager will eventually realize that her utility for choosing the risky project is greater than the safe one. At this point this manager, and hence all others because of their assumed homogeneity, will switch to the risky project. Ideally, the switch would occur at the point where the expected value of the risky project first exceeds that of the safe one. For risk averse managers this will not happen until the expected value of the risky project is much higher than the safe one, thereby leading to a loss to the organization.

With decision insurance however, the switching point is lower and very close to the ideal one that a risk neutral person would choose provided the group size is large. Thus insurance makes the organization have higher expected payoff.

Interestingly, the reverse transition does not take place at the same point when there is insurance. This can be seen by considering that all managers are now involved in the risky project while its success probability decreases over time. At some point, the utility of the safe project will be larger, causing everyone to switch to the safe project. Interestingly however, with insurance in place, this switching point is at a lower value of the probability than it was for the switch from safe to risky. This is due to the larger profit sharing that comes from everyone choosing the risky project as opposed to the safe one. This is because when everyone is choosing a risky project the expected value of profit sharing is higher than if they were using the safe project, and hence they are less risk averse. Thus, the transition for all of them to switch from risky to safe will be at a lower value of the probability of success than the transition the other way. This an interesting social example of hysteresis, which can happen anytime cooperative behavior is at play when constraints change.

4 Discussion

In this paper we proposed a decision insurance mechanism that encourages risk averse decision makers to act as if they are risk neutral. Specifically, we showed how a properly designed decision insurance mechanism can be implemented within organizations so as to increase the likelihood that managers will behave towards the best interest of the enterprise even when they are risk averse. Equally important, we showed how one can use the social network of a decision maker to set up a process that helps reduce the moral hazard inherent in this insurance. This was accomplished by automatically identifying the community of practice around a given manager from patterns in email exchanges in the organization.

From this community of practice one can set up an effective monitoring group whose structure is such that the ability to detect free riding more than compensates the possibility of false positives, ie. construing as free riding what is not on the part of the manager. Likewise, this implementation within the enterprise offers the possibility of addressing adverse selection by making participation mandatory for all managers.

A quantitative analysis of all the factors that enter into this decision insurance identified three possible behavioral regimes. In the first one, members contribute to production but avoid risky projects. In the second regime, managers adopt a risk neutral stance when taking decision insurance. In the third one, they free ride by not giving the full attention to the decision problem at hand while relying on their share of revenue from the efforts of others to compensate for the resulting low performance of their own projects.

An additional advantage of using communities of practice for the purposes of monitoring decision makers is that these communities tend to quickly restructure themselves when new tasks confront the organizations. This is to be contrasted with the very slow response that a formal organization exhibits when having to reorganize in light of new constraints [9].

Implementing this mechanism within an organization would involve the following steps.

1. Identify the participants, ie.. those managers that make decisions and decide whether the decision insurance should be voluntary or applied to every decision maker in the organization.

2. Find the community of practice of those insured to form the monitoring groups for each of the insured. This involves finding the social network and deciding what portion of it to use for the monitoring groups.
3. Set the appropriate threshold values, premiums and payout based on the organizations assessment of the value of the different projects (based on past history, for example).
4. After the period of insurance expires, the organization gets the premiums and pays out the amount of insurance to those whose outcomes failed.

Notice that this mechanism is only useful in situations where the expected return on risky projects is higher than that of safer ones. In addition practical implementation will require that the time horizon for measurable returns be short. Once organizations gain experience with decision insurance over short times they could extend it to longer horizons, perhaps using more sophisticated tools such as futures similar to those in financial markets.

While the particular model we used to quantify the benefit of decision insurance makes simplified assumptions and thus might not describe in detail the workings of a real organization, it nevertheless shows behavioral regimes that can be encountered. Thus it could serve as a guide when trying to implement this mechanism inside an organization. It may be possible to obtain more precise agreement with real organizations by addressing some of its limitations.

As one example, the model uses parameters for some group behaviors that are not readily measured directly. These include the size of the threshold necessary to insure fair detection when free riding with minimal false positives. As one approach to estimating appropriate values, the company could try various pilot systems with differing choices, premiums, groups and thresholds. Observing the resulting participation can give some indication of the appropriate parameters, as well as collecting the necessary information on the track records of managers so that premiums can be set appropriately.

In our analysis we assumed binary choices between a risky and a safe project, rather than the more realistic scenarios of continuous variation of effort. Thus in general it may be the case that only some people in the group are confronted with a moral hazard. The restriction to binary choices highlights this particular dilemma while ignoring the differences among group members.

An important extension of our analysis would incorporate the heterogeneous levels of payoffs, skills and risk attitudes in the group. For instance in a heterogeneous group, adverse selection can be more of a problem because people whose projects involve small payoffs would want to join and free ride on those insured to make decisions that involve lots of money, and therefore higher profit sharing for everyone. This can be accommodated by weighting the premiums and payoffs by the size of the projects.

Another issue is the population dynamics when payoffs change as new projects come into the focus of managers. As illustrated by the model, the population can show interesting hysteresis and time lags and it would be interesting to determine the conditions under which they take place. In particular, to what extent the response to changing constraints depend on the nature of the social network [8]. Furthermore, the natural turnover of organizations imply that newcomers will be joining this insurance scheme. In an insurance scheme it can take a while for newcomers to establish a track record for assessing premiums. One advantage of using the social network structure as part of the mechanism is the ability to gain the needed information in less time than required to complete a number of decisions required to establish a formal track record. This is due to the speed with which an informal network gets established.

Equally interesting will be the experimental test of this mechanism in the laboratory, which will allow for the evaluation of its effectiveness as well as the incentives and tradeoffs that can make it operational inside organizations. Furthermore, a real implementation of decision insurance will expose its value and shortcomings.

And finally, it will be of interest to extend this mechanism outside a formal organizational structure. Outside the enterprise the problem of moral hazard is more severe since there is no formal monitoring structure such as managers. Nevertheless, there are instances of informal groups that are not part of a formal enterprise where peer group monitoring can be sufficient to prevent moral hazard. An example is provided by the microlending institutions in the third world. In some of these cases [6] the group of potential recipient has both the incentive and the knowledge to reduce moral hazard, since the prospect of future loans depends on the repayment history of the group as well as how they use the loan.

Regardless of the simplicity of the model, decision insurance offers a flexible methodology for tackling the vexing problem of making decision makers less risk averse.

References

- [1] S. R. Watson B. Fischhoff and C. Hope. Defining risk. *Policy Science*, 17:123–139, 1984.
- [2] Jerker Denrell and James March. Adaptation as information restriction: the hot stove effect. *Organization Science*, 12:523–538, 2001.
- [3] Kathleen Eisenshardt. Agency theory: and assessment and review. *Academy of Management Review*, 14:57–74, 1989.
- [4] B. Fischhoff, S. Lichtenstein, P. Slovic, S. L. Derby, and R. Keeny. *Acceptable Risk*. Cambridge University Press, Cambridge, 1981.
- [5] David D. Friedman. *Price Theory: An Intermediate Text*. South-Western, 2nd edition, 1990.
- [6] Margaret Grieco. Microfinance and the social exclusion agenda. <http://www.gdrc.org/icm/grieco.html>, 1998.
- [7] Michelle Hanlon, Shivaram Rajgopal, and Terry Shevlin. Large sample evidence on the relation between stock option compensation and risk taking. *mimeo*, 2004.
- [8] Bernardo A. Huberman and Lada A. Adamic. Information dynamics in the networked world. In *Complex Networks*, edited by E. Ben-Naim, H. Frauenfelder and Z. Toroczkai, pages 371,395. Springer, 2004.
- [9] Bernardo A. Huberman and Tad Hogg. Communities of practice: Performance and evolution. *Computational and Mathematical Organization Theory*, 1:73–92, 1995.
- [10] A. F. Herbst J. F. Marshall, V. Bansal and A. L. Tucker. Hedging business cycle risk with macro swaps and option. *Continental Bank Journal of Applied Corporate Finance*, 4:103–108, 1992.
- [11] D. Kahneman and D. Lovallo. Timid choices and bold forecasts: a cognitive perspective on risk taking. *Management Science*, 39:17–31, 1993.
- [12] D. Kahneman and A. Tversky. Prospect theory: An analysis of decision under risk. *Econometrica*, 47:143–157, 1979.
- [13] T. Lant. Aspiration level adaptation: an empirical exploration. *Management Science*, 38:623–644, 1992.

- [14] James March. Variable risk preferences and adaptive aspirations. *Journal of Economic Behavior and Organizations*, 9:5–24, 1988.
- [15] James March. *A Primer on Decision Making*. The Free Press, New York, 1994.
- [16] James March and Zur Shapira. Managerial perspectives on risk and risk taking. *Management Science*, 33:1404–1418, 1987.
- [17] Anthony M. Marino and Jan Zabochnik. Profit centers and incentives in teams. CLEO Research Paper C01-22, USC, May 15 2001.
- [18] M.Douglas. Risk as a forensic resource. *Daedalus*, 119:1–16, 1990.
- [19] Robert C. Merton. *On Consumption-Indexed Public Pension Funds*. National Bureau of Economic Research. University of Chicago Press, 1983.
- [20] Matthew Rabin and Richard H. Thaler. Risk aversion. *Journal of Economic Perspectives*, 15:219–232, 2001.
- [21] Shivaram Rajgopal and Terry J. Shevlin. Empirical evidence on the relation between stock option compensation and risk taking. *J. of Accounting and Economics*, 33(2):145–171, 2002.
- [22] W. Samuelson and R. J. Zeckhauser. Status quo bias in decision making. *Journal of Risk and Uncertainty*, 1:17–59, 1988.
- [23] Zur Shapira. *Risk Taking: a Managerial Perspective*. Russell Sage Foundation, 1995.
- [24] Robert J. Shiller. Expanding the scope of individual risk management: Moral hazard and other behavioral considerations. *Economic Notes*, 26:361–368, 1997.
- [25] Robert J. Shiller. *The New Financial Order: Risk in the 21st Century*. Princeton University Press, 2003.
- [26] P. Slovic. The relative influence of probabilities and payoffs upon perceived risk of a gamble. *Psychonomic Science*, 9:223–224, 1967.
- [27] P. Slovic. Perception of risk. *Science*, 236:280–285, 1987.
- [28] Joshua Tyler, Dennis Wilkinson, and Bernardo A. Huberman. Email as spectroscopy: Automated discovery of community structure within organizations. In *Communities and Technologies*, edited by M. Huysman, E. Wegner and V. Wulf. Kluwer Academic, 2003.

- [29] W. Waggenar. Generation of random sequences by human subjects: a critical review of the literature. *Psychological Bulletin*, 7:65–72, 1972.
- [30] Fang Wu and Bernardo A. Huberman. Finding communities in linear time: a physics approach. *European Physics Journal*, B38:331–338, 2004.