

Solving the Organizational Free Riding Problem with Social Networks

Tad Hogg and Bernardo A. Huberman

HP Labs
Palo Alto, CA

Abstract

We describe how social networks can help reduce certain situations of free riding in organizations. The technique uses the homophily of an organization's social network to select an advisory monitoring peer group likely to share knowledge and skills with each participant, and therefore most able to detect free riding in situations requiring those skills. We illustrate this application in the context of a new mechanism, decision insurance, which helps align decision makers' risk preferences with those of their organization.

Introduction

Social dilemmas often arise when people attempt to produce a common good in the absence of central authority, and can lead to a "tragedy of the commons" (Hardin 1968) with poor outcomes for the group. These dilemmas are also relevant to the adoption of new technologies and the mobilization of political movements.

In a prototypical example of social dilemmas (Axelrod 1984), each person has two choices: to contribute to the common good or to shirk and free ride on the work of others. The relative payoffs of these choices lead to a social dilemma when everyone shares equally in the common good, regardless of their actions, but incurs a cost if they choose to contribute. Each person who cooperates increases the common good a bit, but shares that increase with everyone else. When the cost of cooperating is greater than the marginal benefit, the individual is tempted to defect. Each individual faces the same choice so all defect and the common good is not produced at all. The individually rational strategy of weighing costs against benefits results in an inferior outcome. Such social dilemmas are more likely as group size increases (Olson 1965), since each person's contribution is relatively smaller and their actions are less visible to others in the group.

While social dilemmas arise in a variety of contexts, organizations whose production depends on individual judgments involving significant specialized skill and knowledge face a particularly subtle version. This *organizational social dilemma* for knowledge-intensive work arises when free riding is difficult for most people to recognize because they

lack sufficient knowledge to evaluate observed behavior. However, the people in the organization are heterogeneous. The few people whose knowledge is similar to the free rider are likely to be able to recognize free riding. Furthermore, the public recognition of free riding can be enough to inhibit it, such as through loss reputation among peers or punishment. These characteristics of heterogeneous population and cost associated with revealed free riding contrast with other scenarios involving social dilemmas.

In this paper we describe how social networks facilitate cooperation in this organizational social dilemma. As an application, we then illustrate how this approach addresses the moral hazard arising in organizations using an insurance mechanism to manage risk among the decision makers.

Detecting Free Riding via Social Networks

Social networks can help address organizational social dilemmas because individuals with similar characteristics tend to associate, a phenomenon known as homophily (McPherson, Smith-Lovil, & Cook 2001). Thus probability to detect free riding is higher for those closer in the network due to their similarity of skills and knowledge. This observation is the key rationale for using social networks to help reduce organizational free riding. Specifically, we propose using these networks to identify a "monitoring group" likely to be able to detect free riding. This small peer group can enforce social norms informally, reducing the greater tendency toward free riding in large groups due to increased anonymity. If this is insufficient, reports from the group could also lead to formal investigation and sanctions within the organization.

To model the resulting incentives, during a given evaluation period, we suppose the probability one person detects free riding by another depends only on distance d between them in the network. We denote this probability by p_d and assume independent detections so the chance at least one person among a group of monitors $\{A_1, A_2, \dots\}$ detects free riding is

$$1 - \prod_i (1 - p_{d_i}) \quad (1)$$

where d_i is the network distance between A_i and the evaluated person. Here detection occurs when any single group member reports free riding, but Eq. (1) readily generalizes to requiring multiple reports.

Many observed correlations in behavior and preferences in social networks decrease as a power law. As a specific example, we take p_d to have a power law form:

$$p_d = \alpha d^{-\beta} \quad (2)$$

Here α , between 0 and 1, is the probability an immediate neighbor detects free riding and β characterizes the decreasing probability with distance in the network, so $\beta > 0$, with typical observed values for power laws in network correlations having β ranging between 1 and 3.

In practice, the available social network used to select monitors will not exactly match the actual network in the group. Instead the network data is only an approximation of people's relationships, preferences and knowledge of each other. Thus the network used in the mechanism could be incomplete, noisy or out of date. A significant question for using social networks to reduce free riding is the extent to which the required incentive structure is affected by such lack of knowledge.

As an initial step toward addressing this question, we consider a sample of people with multiple known networks, each with differing link semantics and degree of homophily. Treating one of these networks as correctly indicating the ability to detect free riding, according to Eq. (2), then allows using the other networks, or deliberately degraded versions of the networks, as example noisy networks available to the mechanism. One such sample arises from the networks created by users of Essembly¹, which allows members to vote on various ideological issues. The site offers both a social network (FRIENDS) and two ideological networks (ALLIES and NEMESSES). These networks enable users to distinguish between their friends and people on the site whom they tend to agree (or disagree) with but may not know personally. Users form these networks explicitly and each link must be approved by both users involved in that link.

Our data set consists of fully anonymized network connections (i.e., friends, allies and nemeses) as of December 12, 2006. There are 15,092 unique registered users who voted, and 5,028 of them have declared friends, allies, or nemeses. The networks differ in their links, e.g., less than a third of neighbors in the FRIENDS network are also neighbors in the ALLIES network, and vice versa. Nevertheless, the networks are somewhat correlated in that, say, a friend is more likely to be an ally than is a random other user. For this paper, we use these networks only as examples of networks explicitly created by users, with different nominal meanings for the links. Essembly voting itself does not include the risk and payoff aspects of social dilemmas, and in particular is not a community with free riding. Nevertheless, the Essembly networks can illustrate the effects of using incomplete knowledge of a social network in the context of a model of free riding detection given in Eq. (2).

Specifically, we suppose the social network in Essembly (the FRIENDS network) gives the graph structure defining likelihood of detection, i.e., determining the distance in Eq. (2) for each pair of people. As the organizational group




Figure 1: Detection probability Eq. (1) for the Essembly networks among the top 100 users as a function of β with $\alpha = 0.5$ using monitoring groups of size 5 from among those users. Curves from top to bottom are for monitors determined from full knowledge of the FRIENDS network (black), just half the links in the that network averaged over 100 randomly selected subgraphs (dashed, black), the ALLIES network (dashed, gray) and by random selection (gray).

G we use the 100 most active users in our dataset. We consider four approaches to selecting a monitoring group of size $m = 5$ for each person u :

1. select the m closest people to u in the FRIENDS network
2. select the m closest people to u in the ALLIES network
3. select the m closest people to u in a partially observed version of the FRIENDS network
4. select m people randomly

in all cases selecting the monitors from among the members of the group G other than the person u . These selection methods compare performance without benefit of social network information (random selection), using full knowledge of the relevant network and two types of noise in applying social networks: 1) the available network is correlated with the relevant network, but formed with different intent in the links (in this case reflecting commonality of behavior in terms of voting preferences rather than any necessary implication that linked people personally know each other), and 2) the available information is incomplete because people have not explicitly made their links known to the system.

Fig. 1 compares these methods. As one would expect, using the network whose distances correspond directly to the ability of people to detect free riding gives the best performance. Performance degrades with noise, but the assumed network provides a useful monitoring group as long as its links have reasonable correlation with monitoring ability. Thus, at least for these networks, the properties needed for detecting free riding degrade gracefully with the types of noise likely to occur in available networks.

Managing Risk in Organizations

In this section, we present an example of using social networks to reduce organizational free riding that could be in-

¹Essembly LLC at www.essembly.com

roduced by a new mechanism for managing risk attitudes among decision makers. Due to its emphasis on detailed risk assessments, this mechanism is well-suited to using the homophily properties of social networks.

While most organizations consider taking risks essential for success, few have policies to encourage appropriate risk taking by their employees. Instead organizations often reward outcomes, not decisions, which inhibits risk taking by risk averse decision makers. Moreover, the negative connotation associated with gambling (March 1994; March & Shapira 1987) makes risky choices that fail appear to be mistakes that could be avoided. Additional difficulties arise from how individuals act when confronted with risk (Kahneman & Lovallo 1993), misperceptions of randomness (Waggenar 1972) and aversion to ambiguity (Camerer & Weber 1992). Furthermore, risky alternatives suffer from framing issues: people tend to be risk averse with formulations emphasizing positive outcomes, but risk seeking when emphasizing poor outcomes (Kahneman & Tversky 1979). In practice, managers often focus on the value of the outcome rather than weighing the value by its probability (Shapira 1995; B. Fischhoff & Hope 1984; Douglas 1990). Moreover, managers often do not understand or trust probabilistic measures (Slovic 1967; Fischhoff *et al.* 1981; Slovic 1987), and are often reluctant to change behavior because of status quo biases (Samuelson & Zeckhauser 1988; Denrell & March 2001). This mismatch in risk preferences between individuals and organizations motivates mechanisms to encourage risk averse managers to make decisions based on their assessment of the best interest of the enterprise.

Insurance is a common approach to risk, by spreading the risk over a large group. This mechanism allows risk averse individuals to trade an uncertain but possibly large loss for a certain but smaller loss (the premium). Insurance works well when the number of insured people is large, their risks are not highly correlated and sufficient experience with the risks enables setting appropriate premiums.

Several economists suggest extending insurance to a variety of risky endeavors that are difficult for individuals to insure against. Examples include consumption-indexed public pension funds (Merton 1983), options on business cycle variables such as consumer confidence indices (J. F. Marshall & Tucker 1992), and insurance for the choice of a professional career ("career insurance") and volatile property values (Shiller 2003). Potentially, such insurance could improve economic performance by encouraging individual choices that would benefit society as a whole, on average, but are too risky individually. While provocative, these proposals are difficult to implement because they require finding groups willing to assume these risks and gaining enough experience in assessing risk and setting premiums.

As with other insurance, these proposals must also contend with two key challenges. The first challenge is moral hazard: the tendency to change behavior once insured so the insured outcome becomes more probable. The second challenge is adverse selection: those seeking insurance tend to be those with higher than average risk but this increased risk is not apparent to the insurer.

Decision Insurance

Within organizations, managers should choose projects with the highest expected value for the company. However, when managers believe risky projects have a higher expected value than less risky alternatives, risk averse managers will be reluctant to take the risk when their performance is largely evaluated by success of the projects they select to undertake. Decision insurance allows managers to transfer some of this risk to the organization by paying a premium in exchange for a payoff from the group's profits if the project fails. Decision insurance can adjust for variation in risk attitudes among people by having different combinations of pay for performance and profit sharing for different managers, which amounts to having different insurance premiums. This flexibility, as well as senior managers' ability to decide who participates, helps address adverse selection. Decision insurance can also be viewed as a form of profit sharing, a perspective likely to be more familiar to participants than insurance.

Decision insurance has several aspects:

- **scope:** Identify the relevant decisions and outcomes, and quantify the timing and objective success measures. The outcomes could occur at different times. This process is similar to scoping other information-based mechanisms to aid decision making (Hahn & Tetlock 2006).
- **group:** Identify the participants, i.e., those managers who make the decisions and decide who should participate in the decision insurance, e.g., whether the insurance should be voluntary or a mandatory form of profit sharing for every decision maker in the group.
- **social network:** Find the community of practice for each of those insured, using available social networks within the organization as described below. These communities form the monitoring groups for the insured.
- **parameters:** Set appropriate threshold values, premiums and payouts based on the senior managers' assessment of the value of the different projects.
- **operation:** The organization collects the premiums and pays out the amount of insurance to those whose projects failed at the end of each insurance period.

Setting payouts and premiums depends on the nature of the decisions. For example, if the value of success is widely known (e.g., bidding on a government contract with specified value) but the probability of success is not known (e.g., due to potential bids from competitors), the manager evaluating whether this project is worthwhile for the organization (i.e., has high expected value) would need to devote resources to better estimating the probability of success and evaluating how to increase that probability. In this situation, the known value could form the basis for the insurance payout in case of failure. If, on the other hand, the value to the organization must be estimated as part of the decision process, then senior managers setting insurance payouts would have to use less specific information, e.g., the values of similar projects in the past. The overall scale for the payouts can vary dynamically with the profitability of the organization as a whole. This makes the mechanism suited to organizations




Figure 2: Schematic behavior regimes for selection among a safe project, a risky one and free riding. The dashed curve shows the ideal behavior: switching between safe and risky projects based on which has the higher expected value. Risk averse managers require a larger expected value to overcome the risk when there is no insurance (gray line). Decision insurance shifts preferences closer to the ideal, but could instead lead to free riding if detection probability is too low (black lines). The free riding regime applies only to the case of decision insurance, representing the effects of the moral hazard introduced by the insurance.

performing many independent risky projects, so the relative variance in performance for the organization is smaller than for individuals.

The mechanism's operation must be credible to participants. For example, to encourage more risk taking the insurance payouts on failure should not be just a prelude to other sanctions, such as loss of status or even employment.

Fig. 2 schematically illustrates choices when a manager selects between a safe and a risky project. Ideally, the manager would pick the project with the highest expected value after evaluating the alternatives. Without insurance, risk averse managers tend to favor the safe project, even when they believe its expected value is less than the risky one. Decision insurance improves this situation by having the managers act closer to risk-neutral decision makers, but may also lead to free riding where managers do not give full effort in evaluating their choices. Details of the boundaries between these regimes, which could differ among people in the group, depend on the decision makers' risk attitudes and the payoffs of the various projects.

Moral Hazard and Community of Practice

With decision insurance, moral hazard arises through the effort managers exert in evaluating the success likelihood or value of potential projects. With compensation no longer tied primarily to the success of the individual's efforts, there is a temptation to work less and suffer only a small decrease in individual compensation as a share in the lowered profits of the organization. Even if group members don't actually free ride, the appearance of a conflict of interest caused by insurance could inhibit its widespread adoption.

Free riding is conventionally addressed through monitoring or reputation. Managers in an organization provide some performance monitoring. However especially in the context of knowledge work, senior managers may lack the time or detailed expertise to know whether projects' expected values are accurately estimated through extensive application of a decision-maker's expertise and familiarity with the situation.

To discourage free riding, we propose supplementing manager oversight with a peer group of coworkers likely to be familiar with the decision maker's choices². Specifically each member of the group assesses whether the manager is free riding in the context of a decision making event. If the number of group members reporting free riding exceeds a threshold value, then senior management can decide to investigate and, if appropriate, impose a penalty on the decision maker. This penalty could take the form of an increased decision insurance premium or other organizational sanction.

The monitoring group size and investigation threshold should balance the need to deter free riding against the possibility of false positives, which would tend to make people hesitate to take risky projects. That is, peer monitoring may encourage people to behave as they expect the majority of the group thinks is appropriate. This would counteract the benefit of the insurance when the majority is more risk averse than the insured or group members have significantly less ability to evaluate the decisions. Thus although the monitoring group can reduce free riding, care must be taken that it does not adversely affect the decision making of the insured.

Ideally, the threshold should be large enough to avoid false positives, while still having a reasonable chance of detecting free riding when it occurs. Such a choice ensures the formal investigation is rarely exercised while the group's mere existence and the insured's desire to maintain reputation among the peer group discourages free riding.

Whether this ideal situation is possible depends on the ability of group members to identify free riding. Thus a key aspect of avoiding free riding among decision makers with specialized skills and knowledge is finding a suitable monitoring group. As we describe in the remainder of this section, available social networks provide one approach to identifying a suitable group for each insured decision maker.

Establishing an effective monitoring group requires identifying coworkers familiar with both with the manager and the nature of the decisions. This can be challenging in rapidly changing organizations or projects. Fortunately, those belonging to the manager's community of practice often share both skills and relevant knowledge so are likely to be able to evaluate decision quality as distinct from just observing the eventual outcomes of the decisions. Thus an obvious strategy is to determine those communities of practice from available information on the social network inside

²This group is reminiscent of the so-called friendly societies that were common prior to modern insurance companies and in which people of similar social background provided mutual support while being able to monitor for free riding.

the organization.

The advent of electronic communications offers a unique opportunity to observe the flow of information along both formal and informal channels. For example, email exchange is an indicator of collaboration and knowledge exchange (Huberman & Adamic 2004). 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 (Tyler, Wilkinson, & Huberman 2003; Wu & Huberman 2004; Schwartz & Wood 1993).

The monitoring group can have explicit directions to evaluate the decision-maker. However people in the group, as members of the organization, also have direct incentives to ensure no one free rides. This incentive arises from their profit sharing based on the organization's performance, so they face a loss if others free ride. Thus members of the monitoring group have both incentives for monitoring and the ability to evaluate decisions due to their commonalities with the decision maker.

Moreover, in contexts where free riding may occur, people can prefer to join organizations with monitoring groups to help enforce social norms of productive behavior (Gurker & others 2006). The expectation monitoring encourages greater overall performance can more than compensate for perceived intrusiveness of the groups and the time required to participate in the monitoring groups. Thus, such uses of social networks can encourage positive self-selection and reduce situations in which the monitors would need to act.

Discussion

In this paper we proposed organizations use available social networks to reduce free riding, thereby enabling their use of economic mechanisms that would otherwise be rejected due to fear of free riding. We described one such mechanism: decision insurance. Decision insurance offers a flexible mechanism for tackling the vexing problem of making decision makers less risk averse in situations where the expected return on some risky projects is higher than that of safer ones, but senior managers lack the time or detailed knowledge to evaluate those projects themselves (in which case they could simply specify which projects their organizations should work on).

To reduce free riding incentives, social networks allow more fine-grained monitor selection than just using the closest individuals. For instance, if people are reluctant to critique friends or more likely to collude with them, even if that reduces their profit sharing, the selection of the monitoring group could be restricted to avoid direct neighbors, with perhaps some loss in ability to detect free riding. As another example, social networks allow individually adjusting the mechanism based on position in social network, e.g., lower insurance premium for those in more highly connected regions of the network since they are less likely to avoid detection of free riding. Such adjustment has the further benefit of encouraging people to reveal social network information but also could tempt them to collude with others to reveal links only to those who agree to give them good ratings. Thus using social networks within a mechanism makes most sense

when spoofing the network is relatively costly, e.g., by having the mechanism rely on more distant members of the network than just immediate neighbors, forcing a larger group to collude, and using revealed activities – such as exchanged emails or edits to shared files – that are costly to spoof (i.e., more than just agreeing with someone else to put a link to them in a database). Alternatively, the mechanism could include incentives to reveal the strength of links by having payout partially dependent on how well a person can predict behavior of people they link to, as previously proposed for unstructured groups (Chen, Fine, & Huberman 2004; Prelec 2004).

An advantage of using communities of practice for monitoring decision makers is that these communities tend to restructure more quickly than the formal organization when confronted with new constraints or opportunities (Huberman & Hogg 1995). One common change is turnover in organizations leading to new people continually joining the group of decision makers. Newcomers will require some time to establish a formal track record for accessing premiums, but informal networks for newcomers within the group are often established more rapidly. Thus an interesting question is how the response of the communities depend on the nature of the social network (Huberman & Adamic 2004).

Beyond its use for reducing free riding, the monitoring group could be a proactive aid to the decision maker, as an informally identified group of people who likely have relevant knowledge. Hence the identified group could participate in additional mechanisms of aggregating information (Chen, Fine, & Huberman 2001; Hahn & Tetlock 2006) relevant for the manager.

Using social networks is particularly useful for informal organizations, where a peer group can not be readily estimated from the formal organization structure. These situations include extensive subcontracting or outsourcing where the people involved in any particular formal organization may change rapidly. Thus adding social networks to economic mechanisms could facilitate a broader range of organizational structures.

Our discussion of decision insurance using social networks leads to two empirical questions for future work: how well does insurance alter actual risk behaviors in the context of business decisions, and how to select appropriate parameters to instantiate the mechanism in particular organizations. For instance, laboratory experiments could study controlled decision settings. Initial decision insurance implementations will likely require relatively short times for measurable returns, so people can gain confidence in the mechanism. One approach to estimate appropriate parameters for decision insurance is through various pilot programs with different choices. Observing the resulting participation can give some indication of the appropriate choices, as well as collecting the necessary information on the track records of managers to set premiums. To some extent, this process could be self-organizing as people choose to be part of an organization they feel will likely perform well. Once organizations gain experience with decision insurance, it could be extended to longer time horizons, perhaps using more sophisticated tools such as tradeable futures and options simi-

lar to those in financial markets.

Successful examples within organizations would give confidence and experience to extend such insurance mechanisms more broadly (Shiller 2003). Outside the enterprise moral hazard is more severe since there is no formal monitoring structure such as managers. Nevertheless, informal peer groups, not part of a formal business, can reduce moral hazard. An example is microlending institutions in the third world. In some cases (Grieco 1998) the group of potential loan recipients 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.

Available social networks can be useful in planning targeted marketing via word of mouth (Domingos & Richardson 2001), improving collaborative filtering (Lam 2004), and enhancing online reputation mechanisms (Hogg & Adamic 2004). Reducing moral hazard in decision insurance is a further example of how social networks can improve economic mechanisms. Thus the growing availability of social networks provides numerous opportunities for facilitating a variety of novel economic mechanisms.

Acknowledgements

We thank Michael Brzozowski for analysis of the Essembly data and Chris Chan and Jimmy Kittiyachavalit of Essembly for their help in providing the data. We have also benefited from discussions with Kay-Yut Chen and Murat Kaya.

References

- Axelrod, R. 1984. *The Evolution of Cooperation*. NY: Basic Books.
- B. Fischhoff, S. R. W., and Hope, C. 1984. Defining risk. *Policy Science* 17:123–139.
- Camerer, C., and Weber, M. 1992. Recent developments in modeling preferences: Uncertainty and ambiguity. *J. of Risk and Uncertainty* 5:325–370.
- Chen, K.-Y.; Fine, L. R.; and Huberman, B. A. 2001. Forecasting uncertain events with small groups. In *Proc. of the 3rd ACM Conference on E-commerce*, 58–64.
- Chen, K.-Y.; Fine, L. R.; and Huberman, B. A. 2004. Eliminating public knowledge biases in information-aggregation mechanisms. *Management Science* 50:983–994.
- Denrell, J., and March, J. 2001. Adaptation as information restriction: the hot stove effect. *Organization Science* 12:523–538.
- Domingos, P., and Richardson, M. 2001. Mining the network value of customers. In *Proc. of the 7th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, 57–66. ACM Press.
- Douglas, M. 1990. Risk as a forensic resource. *Daedalus* 119:1–16.
- Fischhoff, B.; Lichtenstein, S.; Slovic, P.; Derby, S. L.; and Keeny, R. 1981. *Acceptable Risk*. Cambridge: Cambridge University Press.
- Grieco, M. 1998. Microfinance and the social exclusion agenda. In *Accounting SIG*. Buisness School, Univ. of North London.
- Gurerk, O., et al. 2006. The competitive advantage of sanctioning institutions. *Science* 312:108–111.
- Hahn, R. W., and Tetlock, P. C., eds. 2006. *Information Markets: A New Way of Making Decisions*. Washington DC: AEI Press.
- Hardin, G. 1968. The tragedy of the commons. *Science* 162:1243.
- Hogg, T., and Adamic, L. 2004. Enhancing reputation mechanisms via online social networks. In *Proc. of the 5th ACM Conference on Electronic Commerce (EC'04)*, 236–237. ACM Press.
- Huberman, B. A., and Adamic, L. A. 2004. Information dynamics in the networked world. In Ben-Naim, E.; Frauenfelder, H.; and Toroczkai, Z., eds., *Complex Networks*. Springer. 371–395.
- Huberman, B. A., and Hogg, T. 1995. Communities of practice: Performance and evolution. *Computational and Mathematical Organization Theory* 1:73–92.
- J. F. Marshall, V. Bansal, A. F. H., and Tucker, A. L. 1992. Hedging business cycle risk with macro swaps and option. *Continental Bank Journal of Applied Corporate Finance* 4:103–108.
- Kahneman, D., and Lovallo, D. 1993. Timid choices and bold forecasts: a cognitive perspective on risk taking. *Management Science* 39:17–31.
- Kahneman, D., and Tversky, A. 1979. Prospect theory: An analysis of decision under risk. *Econometrica* 47(2):263–292.
- Lam, C. 2004. SNACK: incorporating social network information in automated collaborative filtering. In *Proc. of the 5th ACM Conference on Electronic Commerce (EC'04)*, 254–255. ACM Press.
- March, J., and Shapira, Z. 1987. Managerial perspectives on risk and risk taking. *Management Science* 33:1404–1418.
- March, J. 1994. *A Primer on Decision Making*. New York: The Free Press.
- McPherson, M.; Smith-Lovil, L.; and Cook, J. M. 2001. Birds of a feather: Homophily in social networks. *Annual Review of Sociology* 27:415–444.
- Merton, R. C. 1983. *On Consumption-Indexed Public Pension Funds*. National Bureau of Economic Research. University of Chicago Press.
- Olson, M. 1965. *The Logic of Collective Action: Public Goods and the Theory of Groups*. Harvard Univ. Press.
- Prelec, D. 2004. A Bayesian truth serum for subjective data. *Science* 306:462–466.
- Samuelson, W., and Zeckhauser, R. J. 1988. Status quo bias in decision making. *Journal of Risk and Uncertainty* 1:17–59.
- Schwartz, M. F., and Wood, D. C. M. 1993. Discovering shared interests among people using graph analysis. *Communications of the ACM* 36(8):78–89.
- Shapira, Z. 1995. *Risk Taking: a Managerial Perspective*. Russell Sage Foundation.
- Shiller, R. J. 2003. *The New Financial Order: Risk in the 21st Century*. Princeton University Press.
- Slovic, P. 1967. The relative influence of probabilities and payoffs upon perceived risk of a gamble. *Psychonomic Science* 9:223–224.
- Slovic, P. 1987. Perception of risk. *Science* 236:280–285.
- Tyler, J.; Wilkinson, D.; and Huberman, B. A. 2003. Email as spectroscopy: Automated discovery of community structure within organizations. In Huysman, M.; Wegner, E.; and Wulf, V., eds., *Communities and Technologies*. Kluwer Academic.
- Waggenar, W. 1972. Generation of random sequences by human subjects: a critical review of the literature. *Psychological Bulletin* 7:65–72.
- Wu, F., and Huberman, B. A. 2004. Finding communities in linear time: a physics approach. *European Physics Journal B* 38:331–338.