6
Foul Play in Information Markets

Robin Hanson

People have long noticed that speculative markets, though created for other purposes, also do a great job of aggregating relevant information. In fact, it is hard to find information not embodied by such market prices. This is, in part, because anyone who finds such neglected information can profit by trading on it, thereby reducing the neglect.[1]

So far, speculative markets have done well in every known head-to-head field comparison with other forecasting institutions. Orange juice futures improved on National Weather Service forecasts,[2] horse race markets beat horse race experts,[3] Oscar markets beat columnist forecasts,[4] gas-demand markets beat gas-demand experts,[5] stock markets beat the official NASA panel at fingering the guilty company in the Challenger accident,[6] election markets beat national opinion polls,[7] and corporate sales markets beat official corporate forecasts.[8]

Recently, some have considered creating new markets specifically to take

---

For their comments, I thank Robert Hahn and the participants of the 2004 AEI–Brookings conference on information markets. I thank the Center for Study of Public Choice, the Mercatus Center, and the International Foundation for Research in Experimental Economics for financial support.

[1] Andrew Lo, “Finance: A Selective Survey,” *Journal of the American Statistical Association* 95, no. 45 (2000): 629–35; and Koleman S. Strumpf and Paul W. Rhode, “Historical Presidential Betting Markets,” *Journal of Economic Perspectives* 18, no. 2 (2004): 127–41.

[2] Richard Roll, “Orange Juice and Weather,” *American Economic Review* 74, no. 5 (1984): 861–80.

[3] Stephen Figlewski, “Subjective Information and Market Efficiency in a Betting Market,” *Journal of Political Economy* 87, no. 1 (1979): 75–88.

[4] David M. Pennock, C. Lee Giles, and Finn A. Nielsen, “The Real Power of Artificial Markets.,” *Science* 291 (2001): 987–88.

[5] Jakab Spencer, “New ICAP–Nymex Derivatives Have U.S. Gas Markets Number,” *Wall Street Journal*, August 4, 2004.

[6] Michael T. Maloney and J. Harold Mulherin, “The Complexity of Price Discovery in an Efficient Market: The Stock Market Reaction to the Challenger Crash,” *Journal of Corporate Finance* 9, no. 4 (2003): 453–79.

[7] Joyce Berg, Forrest Nelson, and Thomas Rietz, “Accuracy and Forecast Standard Error of Prediction Markets,” Technical Report, University of Iowa, College of Business Administration. 2001

[8] Kay-Yut Chen and Charles R. Plott, “Prediction Markets and Information Aggregation Mechanisms: Experiments and Application” (technical report, California Institute of Technology, 1998).

advantage of these effects. Called prediction markets,[9] information markets, virtual stock markets,[10] artificial markets,[11] or idea futures,[12] such markets are now beginning to estimate such things as product sales, project completion dates, and election outcomes.

Many observers have expressed concerns that such markets might induce various forms of foul play. For example, during the recent furor over the Policy Analysis Market (PAM) of the Defense Advance Research Projects Agency (DARPA), otherwise known as terrorism futures, critics complained that PAM might have allowed bets on the details of individual terrorist attacks.[13] In particular, critics feared that bad guys might do more bad things in order to win bets about those bad events, or might intentionally lose bets in order to reduce market information.[14]

In addition, others have expressed concerns that such markets might induce more people to lie, that markets inside organizations might misdirect resources, perhaps maliciously, and that threats of retribution might limit their effectiveness. In this chapter, I review what we know about these possible forms of foul play, and I suggest some new approaches to dealing with them.

---

[9] Berg, Nelson, and Rietz, "Accuracy and Forecast Standard Error of Prediction Markets"; Justin Wolfers and Eric Zitzewitz, "Prediction Markets," *The Journal of Economic Perspectives* 18, no. 2 (2004): 107–126.

[10] Martin Spann and Bernd Skiera, "Internet-Based Virtual Stock Markets for Business Forecasting," *Management Science* 49, no. 10 (2003): 1310–26.

[11] Pennock, Giles, and Nielsen, "Real Power of Artificial Markets," 987–88.

[12] Robin Hanson, "Market-Based Foresight—A Proposal," *Foresight Update* 10 (1990): 1, 3, 4; Robin Hanson, "Could Gambling Save Science? Encouraging an Honest Consensus," *Social Epistemology* 9, no. 1 (1995): 333; Robin Hanson, "Idea Futures," *Wired* 3, no. 9 (1995): 125.

[13] Actually, PAM would have focused on aggregate geopolitical trends, such as how the chances of political unrest in Saudi Arabia depend on whether U.S. troops leave there. See Charles Polk, Robin Hanson, John Ledyard, and Takashi Ishikida, "The Policy Analysis Market: An Electronic Commerce Application of a Combinatorial Information Market," *Proceedings of the ACM Conference on Electronic Commerce* (Association for Computing Machinery, New York, 2003): 272–73.

[14] Senators, reporters, and economists complained. For example, U.S. senators Ron Wyden and Bryon Dorgan said, "Terrorists themselves could drive up the market for an event they are planning and profit from an attack, or even make false bets to mislead intelligence authorities" (R. Wyden and B. Dorgan, press release, July 28, 2003). Steven Pearlstein argued that "would-be assassins and terrorists could easily use disinformation and clever trading strategies to profit from their planned misdeeds while distracting attention from their real target" (Steven Pearlstein, "Misplacing Trust in the Markets," *Washington Post*, July 30, 2003, E1). Joseph Stiglitz said that "trading . . . could be subject to manipulation, particularly if the market has few participants providing a false sense of security or . . . alarm. . . . The lack of intellectual foundation or a firm grasp of economic principles—or the pursuit of other agendas—has led to a proposal that almost seems a mockery of itself" (Joseph Stiglitz, "Terrorism: There's No Futures in It," *Los Angeles Times*, July 31, 2003.)

Evaluation Standards

Let us begin by considering what standard we should apply when evaluating possible foul play. Information markets are created so that their price estimates can inform the policy choices of a for-profit, nonprofit, or government organization. Those creating such markets would naturally seek the most accurate possible prices on the most valuable topics at the lowest possible costs in terms of time, money, and other relevant resources, including the disruption to existing cultures and practices.

The existence of some estimate error or some resource cost, however, should not by itself be much of a criticism. The relevant benchmark should not be an infeasible perfection, but the accuracy and costs of other social institutions that perform a similar function with similar inputs. In particular, the most relevant benchmark is set by currently used forecasting institutions, such as in-house experts, ad hoc expert committees, independent forecasting agencies, and opinion polls.

Regarding each possible form of foul play, a key question is the degree to which a similar form occurs in competing social institutions. Do information markets pose special concerns?

Limiting Participation

Before discussing specific forms of foul play, let us consider a generic mitigation strategy: limiting participation. Foul play is always committed by someone. Thus, we might hope to limit it in any social institution by limiting who can participate in that institution.

The effectiveness of this strategy depends on our ability either to find indicators of who will engage in foul play or to detect acts of foul play directly. Effectiveness also depends on the rate at which institutional performance degrades with reduced participation.

Information markets can be used to aggregate information from any given set of participants. (Even having just one participant can work.[15]) Although most of the impressive track record of speculative market accuracy has come from markets that are open to all participants, some recent successes, such as some internal markets at Hewlett-Packard, have also come from markets with closed participation.[16] Even so, to the extent that information markets seem especially attractive due to their ability to allow and take advantage of wider participation than other institutions, a special concern must be the additional opportunities for foul play that wider participation might induce.

Let us now review the various forms of foul play: lying, manipulation, sabotage, embezzlement, and retribution.

Lying

[15] Robin Hanson, "Combinatorial Information Market Design," *Information Systems Frontier* 5, no. 1 (2003): 105–119.
[16] Chen and Plott, "Prediction Markets and Information Aggregation Mechanisms."

The participants of a forecasting institution are those who make direct contributions to the forecast. There are also advisors who offer to make suggestions to the participants, and participants may advise other participants. Perhaps the simplest form of foul play is that committed by advisors who lie to, or mislead, the participants. Stock analysts, for example, are often accused of being paid by certain companies to give overly optimistic advice about their company stock.

Misleading advice is a familiar issue with all known existing social institutions, including speculative markets.[17] In general, participants should think skeptically about advice, taking into account any clues they have about advisor track records and incentives. The institutional process that combines participant contributions into forecasts seems largely irrelevant to this skeptical evaluation of advice. If so, all forecasting institutions should have similar levels of this sort of foul play.

Forecasting institutions, however, might vary in the clues they offer to detect lies, or in the incentives they give advisors. If, without information markets, one could get good advice from people who are trustworthy because of their neutral or transparent interests, allowing such advisors to trade could be a problem if doing so would obscure their interests. Advisors who could acquire hidden trading positions might become less trustworthy.

A sufficient way of addressing this problem would be to allow traders to reveal credibly their relevant holdings to those whom they advise. Advisors who chose not to reveal, and not to arrange for neutral holdings, would be treated more skeptically. Revelation should include not only an advisor's direct asset holdings but also any strong shared interests with others who may have such holdings. This is a familiar approach to dealing with conflicts of interests in other areas. Secret accounts, trader anonymity, and complex shared interests, however, might conspire to make it difficult for advisors to reveal credibly their relevant holdings.

Another approach is to expect advisors to reveal their relevant information directly, as participants, rather than indirectly, as advisors. Trading might well be a more effective way for them to reveal their information. If neither of these approaches were sufficient, one might prohibit certain groups of potential advisors from trading.

Manipulation

Another possible form of foul play occurs when participants who want to influence policy decisions directly distort their contributions to the institution forecasts. This is a familiar issue with existing institutions, such as ad hoc expert committees and supposedly independent forecasting agencies. In information markets, this foul play would take the form of people making trades that lose money in order to change prices and hence policy. Even if such trades lost money on average, those losses might be outweighed by gains from more favorable policy.

---

[17] Franklin Allen and Douglas Gale, "Stock-Price Manipulation," *Review of Financial Studies* 5, no. 3 (1992): 503–29.

Even if successful, this sort of manipulation would mainly just reduce the accuracy of information-market prices.[18] As long as decision-makers knew roughly the level of forecast error in prices, such prices could be still useful inputs into decisions. And if such manipulators could similarly bias other forecasting institutions, this would not be a special concern about information markets.

Information markets seem especially hard to manipulate, however. We know of only one apparently successful attempt,[19] and many people have reported failed attempts to manipulate speculative market prices with trades historically,[20] in the field,[21] and in the laboratory.[22] A recent review article concludes that none of these attempts at manipulation had much of a discernible effect on prices, except during a short transition phase.[23] How can this be?

The key thing to understand is that all known speculative markets have a lot of "noise trades," that is, trades made because of mental mistakes, for insurance purposes, or for other noninformational reasons. Furthermore, a manipulator is just another kind of noise trader.

In theory, perfectly rational informed traders could use a subsidized market to aggregate their information and exactly reveal their common estimate.[24] Real markets, however, are full of fools, hedgers, and others whose trades are prompted by factors other than the information they hold. In fact, the opportunity to trade, and win, against noise traders is usually the main profit incentive informed traders have to participate.

If we hold other trading behavior constant, adding more noise trading must increase price errors. But when other traders expect more noise trading, they change their behavior in two important ways.

First, they eagerly scale up the amount they trade for any given amount of information they might hold, as this increases their expected profits. In the limit where the amounts traded are small compared to traders' aggregate risk tolerance, this should fully compensate for the increased noise, leaving the price error exactly the same. That is, as long as there are a few participants with deep enough pockets, or enough participants with shallow pockets, there will be enough people willing to accept the noise traders on

---

[18] Other sorts of manipulation are less of a concern, such as strategic contrary trading by an informed trader to control the rate at which his information is revealed. See Archishman Chakrabortya and Bilge Ylmaz, "Manipulation in Market Order Models," *Journal of Financial Markets* 7, no. 2 (2004): 187–206.
[19] Jan Hansen, Carsten Schmidt, and Martin Strobel, "Manipulation in Political Stock Markets—Preconditions and Evidence," *Applied Economics Letters* 11, no. 7 (2004): 459–63.
[20] Koleman S. Strumpf and Paul W. Rhode, "Historical Presidential Betting Markets," *The Journal of Economic Perspectives* 18, no. 2 (2004): 127–41.
[21] Colin Camerer, "Can Asset Markets Be Manipulated? A Field Experiment with Racetrack Betting," *Journal of Political Economy* 106 (1998): 457–82.
[22] Robin Hanson, Ryan Oprea, and David Porter, "Information Aggregation and Manipulation in an Experimental Market," *Journal of Economic Behavior and Organization* (forthcoming, 2006).
[23] Wolfers and Zitzewitz, "Prediction Markets," 107–26.
[24] Hanson, "Combinatorial Information Market Design," 105–19.

average losing bets.[25]

Now, it may well be true that financial market traders do not fully correct for increases in aggregate noise trading in the world economy, at least along the handful of dimensions that command risk premia. Irrational traders who underestimate the risk they are taking on can create aggregate risks that rational traders cannot afford to eliminate.[26] But this does not seem very relevant for most information markets, which do not estimate aggregate risks.

The second change in behavior is that the increased profit opportunity from more noise traders increases the effort by other traders to obtain relevant information. So, on net, more noise trading should *increase* price accuracy.[27] And, in fact, empirically it seems that financial and information markets with more noise trading, and hence a larger trading volume, tend to be more accurate, all else being equal.[28]

Models of financial market microstructure have considered several types of noise traders, including fools who act randomly, traders with immediate liquidity needs, traders who seek to manipulate a closing price in order to raise their futures-market settlement,[29] and, more generally, traders with quadratic preferences over the market price.[30]

These models verify that manipulators are just another kind of noise trader. A manipulator has hidden information about his bias—that is, how much and in what direction he wants to bias the price. (This includes the possibility of zero bias—that is, of not being a manipulator.) Other traders can respond only to the average expected bias. When the hidden bias is exactly equal to the average bias, it is as if there were no manipulator. When the bias is higher or lower than expected, the price will be higher or lower than expected. But competition among speculators ensures that on average the price is right, and that average price error is reduced when manipulators have larger biases.

---

[25] Note that if aggregate noise trading risk is a problem, then increases in potential participation can reduce this sort of foul-play problem by increasing total pocket depth relative to noise trading.

[26] J. Bradford De Long, Andrei Shleifer, Lawrence H. Summers, and Robert J. Waldmann, "Noise Trader Risk in Financial Markets," *Journal of Political Economy* 98, no. 4 (1990): 703–38.

[27] Albert S. Kyle, "Informed Speculation with Imperfect Competition," *Review of Economic Studies* 56, no. 3 (1989): 317–55; Matthew Spiegel and Avanidhar Subrahmanyam, "Informed Speculation and Hedging in a Noncompetitive Securities Market," *Review of Financial Studies* 5, no. 2 (1992): 307–29.

[28] Joyce Berg, Robert Forsythe, and Thomas Rietz, "What Makes Markets Predict Well? Evidence from the Iowa Electronic Markets," in *Understanding Strategic Interaction: Essays in Honor of Reinhard Selten*, ed. Wulf Albers, Werner Guth, Peter Hammerstein, Benny Moldovanu, Eric van Damme, Martin Strobel, and Reinhard Selten (New York: Springer, 1996), 444–63.

[29] Praveen Kumar and Duane J. Seppi, "Futures Manipulation with Cash Settlement," *Journal of Finance* 47, no. 4 (1992): 1485–1502; Pierre Hillion and Matti Suominen, "The Manipulation of Closing Prices," *Journal of Financial Markets* 7, no. 4 (2004): 351–75.

[30] Hanson, Oprea, and Porter, "Information Aggregation and Manipulation."

Even if manipulators reduce price error on average, however, they might still increase the harm from price errors. Imagine that the harm from a price error depended not just on the magnitude of the error, but also on some additional state that was positively correlated with the hidden manipulator bias. For example, in a market estimating the chance of a terrorist attack, terrorists might perhaps arrange for the size of the attack to be correlated with the forecast error. The market might then become more accurate in estimating whether an attack would occur, but it would also miss the big attacks more often. In such a case, the expected harm from price errors could increase with more manipulation, even as the expected error decreased.

One approach to mitigating this problem is via the parameters that markets estimate. The closer those parameters are to the actual decision parameters of interest, the less likely should be the existence of hidden states that modulate the magnitude of the harm from estimation errors and that are correlated with some manipulator bias. For example, it would be better for a terrorist-attack market to estimate the harm caused by the attack and not just whether an attack occurs.

Sabotage

A potentially very serious form of foul play is where people cause harm to gain trading profits. For example, some people suspected that the September 11, 2001, terrorist attacks on the New York World Trade Center were funded in part by trades of airline stock options. Similarly, some feared that the 1982 Tylenol poisonings were done to profit from short sales on the Tylenol stock. Airline stock prices did fall on September 11, as did the Tylenol stock with the 1982 poisonings. And a study has found that Israeli stock and currency prices respond to Israeli suicide bombings.[31] Thus, it is not crazy to think that terrorists might use financial markets to profit from their acts. Nevertheless, we know of no examples of anyone using financial markets to profit from such sabotage. A thorough study of the September 11 attacks found nothing suspicious, and no trades were ever linked to the Tylenol poisonings.[32] The closest example I can find is the case of Roger Duronio, a well-paid PaineWebber employee who, in 2002, set off a logic bomb in one thousand company computers after investing $20,000 in options betting that the stock price would fall. The damage totaled $3 million, but system redundancy prevented any loss of data, the stock price did not fall, and Duronio was soon caught.[33]

We do, however, know of examples of murder committed to gain life insurance, where the insurance was purchased with this plan in mind. Thus, speculation on sabotage is possible when one person can acquire a large enough stake in an asset whose value he can influence directly enough. We also have examples of extortion of large corporations

---

[31] Rafi Eldor and Rafi Melnick, “Financial Markets and Terrorism,” *European Journal of Political Economy* 20, no. 2 (2004): 367–86.
[32] Thomas H. Kean, Lee H. Hamilton, Richard Ben-Veniste, Bob Kerrey, Fred F. Fielding, John F. Lehman, Jamie S. Gorelick, Timothy J. Roemer, Slade Gorton, and James R. Thompson, *The 9-11 Commission Report*, http://www.9-11commission.gov/report/911Report.pdf (accessed July 2004).
[33] Andy Geller, “Pained Webber: Geek Tried to Sink Stock with Cyber Bomb,” *New York Post*, December 18, 2002.

by people who first demonstrate their ability to cause large amounts of damage. Compared to speculating on sabotage, the extortion strategy runs a greater risk of detection but requires less capital to implement.

A simple interpretation of these facts is that the need for secrecy makes it very hard for skilled labor and willing capital to find each other to implement the strategy of speculating on sabotage. Because relevant prices move for other reasons, one needs a large portfolio of sabotage acts to be reasonably confident of a net profit. But those who are well-positioned to commit a single act of sabotage are usually not well-positioned to commit a stream of such acts. A willing source of capital would thus have to find many skilled saboteurs and would risk detection with each new potential saboteur contacted.

Information markets are typically very thin compared with most financial markets, with relatively little money changing hands. All else being equal, this makes them poor places to speculate on sabotage. Nonetheless, financial markets are also typically tied to large economic aggregates, which are difficult for individuals to influence reliably. If information markets are created to estimate smaller-scale social processes that individuals could more directly influence, speculating on sabotage might be more of an issue. For example, a company might create a market on whether a certain project will meet its deadline, and many individual employees might have the ability to sabotage the project and delay its completion.[34]

One can try to deal with this problem by only estimating large aggregates, by limiting participation,[35] or by allowing investigators of suspicious events to see who made what trades.[36] Another approach, however, is to set limits on trading positions. For each class of traders, one might limit how far asset holdings could move in dangerous directions. For example, regarding the corporate project completion market, the company might estimate bounds on the current implicit stake and minimum acceptable stake for different classes of people. Each employee working on the project might be expected to gain at least $200 worth in professional reputation should the project be completed on time, while a benefit of $100 would be considered sufficient to ensure that he did not harm the project.

Given these assumptions, such an employee could safely trade until he reached a position where he would gain $100 via bets if the project were not completed on time. That is, if he started with no bets, he could be allowed to pay $50 for the asset, “Pays $100 if project misses deadline.” At that point he would still be set to gain at least $100 if the project were completed on time, and so he would not be tempted to sabotage the

---

[34] A perhaps bigger problem with such project deadline markets is that managers often manipulate employee overconfidence and uncertainty about deadlines in order to get more work out of them. See Magne Jorgensen, Karl H. Teigen, and Kjetil Molokken, “Better Sure Than Safe? Overconfidence in Judgment Based Software Development Effort Prediction,” *Journal of Systems and Software* 70, no. 1–2 (2004): 79–83. Such managers would resist the introduction of these markets.

[35] Limiting participation for this reason is analogous to having regulations requiring that one have an “insurable interest” to buy insurance.

[36] Tad Hogg and Bernardo A. Huberman, “Avoiding Moral Hazards in Organizational Forecasting,” Hewlett-Packard Labs, http://www.hpl.hp.com/research/idl/papers/moral/moral.pdf (accessed June 2004).

project.

What if some employees were already at their asset limits, such as having a zero initial stake, where any negative stake is considered dangerous? Well, being that close to a dangerous boundary seems unwise. The company should probably give them all an explicit bonus contingent on project completion and then allow them to trade this down to zero or up to infinity. To make this approach work, one might have to worry about whether people could trade via multiple accounts and whether they had shared interests with others whose stakes should be limited.

Embezzlement

Many observers are concerned that information markets inside organizations could misdirect time, money, and credit, perhaps maliciously. If one creates real-money markets on company-related events, where employees can bet large sums, they may shirk on other tasks to play the market. But if one creates play-money markets, or real-money markets where only small sums can be bet, they may not see why they should bother to participate. How can markets induce enough, but not too much, effort?

Those who choose market topics might do so in part to reward their friends. For example, creating a market on future sales might reward those who have first access to relevant organizational data on sales. Also, team members may withhold insights from team production in order to gain more cash in the markets privately.

These sorts of difficulties with creating explicit monetary reward schemes are ubiquitous in most organizations. Consequently, most employees are not given direct financial rewards on most tasks. Instead, they are usually rewarded on the basis of overall performance evaluations. Such evaluations consider many relevant indicators, but usually no commitment is made to any particular formula for combining those indicators. This allows managers more flexibility to notice and correct for the sort of foul play that direct and formulaic monetary rewards might induce.

A similar performance evaluation approach can be used to deal with foul play in information markets. A standard salient entry in an employee evaluation is that the employee, alone or with some group, initiated a change that was estimated to have added so many dollars to the organization’s bottom line. I suggest that we design internal organization information markets to facilitate similar statements by introducing a new color of money. Creating different colors of money, with limits on their convertibility, is a standard accounting technique for dealing with complex incentive problems in organizations.

Instead of betting cash or play-money, we could enable bets of information-money. Initial holdings of info-money would be distributed not only to groups and individuals, but also to automated market-makers on chosen trading topics. Some specialists would estimate the value of more accurate information on each trading topic, and then each topic’s market-makers would be subsidized at a level corresponding to the estimated value of information on that topic.[37]

Given such subsidies, traders would, on average, increase their holdings of info-money as they made market prices more accurate, and the total increased holdings would

[37] Hanson, “Combinatorial Information Market Design,” 105–19.

correspond to the estimated total value of the information produced. Employees or groups who could show they had consistently increased their info-money holdings could then claim credit on their evaluations. They would claim the amount of their increased holdings as a dollar-valued contribution to the organization's bottom line. This approach would give managers a reasonable basis for allocating the efforts of their subordinates among various tasks, including various info-production tasks.

Statistical analyses of the history of each person's or group's trades would be needed to distinguish consistent increases from mere random fluctuations. Those people with consistent decreases could be encouraged to change something or stop. Those with large but inconsistent fluctuations could be encouraged to keep their fluctuations small until they learned how to make consistent contributions. And those who were afraid to make any trades for fear of losses could be encouraged to make only small trades, where losses need not be stigmatized, while they learned how to make consistent contributions.

To discourage individuals from embezzling team information, one might have the team account trade first on any new team information and only afterward allow individual team members to trade on their own accounts. This could allow individual team members to write dissenting minority reports while avoiding embezzlement of team information.

In cases where there is concern that members might withhold their insights from the team, the team might be given the right of first refusal on member trades, so that a trade would be an individual trade only if the team did not want to make it as a team trade. This approach, however, would make it difficult to allow for anonymous dissenting opinions.

One would want to avoid wasteful contests by different groups to be the first to arrive at the market with easily collected information that is not especially time-critical. This might be done by creating standard processes that trade on such information. Only after this standard trading were done would one let others trade on the information, to express any different beliefs they might have about how exactly such information should be incorporated. One might also slowly raise the subsidy level on a topic from zero, to entice the cheapest possible info supplier to supply it first.

Retribution

Existing forecasts are often inaccurate because someone wants them to be so. For example, a salesman may want to create low expectations about future sales so that his efforts will look good by comparison. Or someone proposing a new project may want to create high expectations so that his project will be approved. Such people often distort the information they present to others in order to create inaccurate forecasts.

If only a few insiders knew that these forecasts were inaccurate, and if information markets threatened to entice those insiders to rat on the deceptive forecasts, then those who preferred the deception might threaten retribution against anyone who contradicted them in the markets. For example, a project leader might punish anyone on his project team who disputed his rosy forecast.

Because similar processes exist in other forecasting institutions, this approach does not appear to create a special concern for information markets. In fact, information markets can substantially reduce this problem via anonymous trading. Anonymous trades can avoid retribution, at least if the leader is not willing to punish all team members whenever the market goes against him. This approach can, however, require a lot of routine anonymous trading to take place so that the mere fact that one is trading anonymously does not make one a target of retribution. Anonymity can also conflict with giving teams a right of first refusal on team member trades.

Anonymous trading can be consistent with allowing managers to oversee their employees' allocation of effort. Even if managers are not able to see individual trades, they might see the time that their subordinates spend trading, along with statistics on their overall trading performance.

Conclusion

The impressive accuracy of information markets, relative to competing forecasting institutions, is encouraging their wider application, but many people have expressed concerns that such markets might encourage various forms of foul play, including lying, manipulation, sabotage, embezzlement, and retribution. I have reviewed each of these forms and provided strategies for mitigating them.

The standard for evaluation should be how information markets compare to competing forecasting institutions, and limiting participation is a generic but crude strategy for limiting foul play.

Inducing lies is only a special concern of information markets when such markets have wider participation than other institutions. Reasonable solutions include having advisors trading instead of talking, or giving them the ability to show their neutral trading position.

Manipulation seems a much weaker concern for information markets than for competing institutions, as manipulative trading should usually improve price accuracy. Manipulation should be a potential problem only when all traders are very risk-averse, or when the harm from price errors correlates in unusual ways with those errors.

Sabotage is not a concern when markets estimate large social aggregates that are hard for individuals to influence, or when the trading stakes are too small to pay for any substantial sabotage efforts. When the events are small enough relative to the trading stakes for sabotage to be a concern, one can limit participation, reveal trades to investigators, or place bounds on individual trading stakes.

To discourage the embezzlement of time, money, and credit within organizations, internal markets could trade a new color of money. When trading topics are subsidized at their value of information, those with consistent trading gains can take credit for adding so many dollars to the organization's bottom line. Standard information sources should have special processes that trade on them before others, and teams should trade on team information before team members do.

Retribution does not seem a special concern of information markets, and anonymous trading can greatly reduce the ability to suppress information through threats of retribution.

Overall, none of these forms of foul play seems worse in information markets when one holds constant who can participate in the forecasting institutions. Allowing information markets to have broader participation than other institutions can bring in more people who may engage in foul play, but many approaches are available for limiting this problem to tolerable levels.