University of Chicago Law School

Chicago Unbound

Public Law and Legal Theory Working Papers

Working Papers

2007

Deliberating Groups versus Prediction Markets (or Hayek's Challenge to Habermas)

Cass R. Sunstein

Follow this and additional works at: https://chicagounbound.uchicago.edu/public_law_and_legal_theory

Part of the **Law Commons**

Chicago Unbound includes both works in progress and final versions of articles. Please be aware that a more recent version of this article may be available on Chicago Unbound, SSRN or elsewhere.

Recommended Citation

Cass R. Sunstein, "Deliberating Groups versus Prediction Markets (or Hayek's Challenge to Habermas)" (University of Chicago Public Law & Legal Theory Working Paper No. 146, 2007).

This Working Paper is brought to you for free and open access by the Working Papers at Chicago Unbound. It has been accepted for inclusion in Public Law and Legal Theory Working Papers by an authorized administrator of Chicago Unbound. For more information, please contact unbound@law.uchicago.edu.

CHICAGO

PUBLIC LAW AND LEGAL THEORY WORKING PAPER NO. 146

DELIBERATING GROUPS VERSUS PREDICTION MARKETS (OR HAYEK'S CHALLENGE TO HABERMAS)

Cass R. Sunstein

THE LAW SCHOOL
THE UNIVERSITY OF CHICAGO

January 2007

This paper can be downloaded without charge at the Public Law and Legal Theory Working
Paper Series: <http://www.law.uchicago.edu/academics/publiclaw/index.html> and
The Social Science Research Network Electronic Paper Collection:
http://ssrn.com/abstract_id=956189

Deliberating Groups versus Prediction Markets (or Hayek's Challenge to Habermas)

Cass R. Sunstein^{*}

Abstract

For multiple reasons, deliberating groups often converge on falsehood rather than truth. Individual errors may be amplified rather than cured. Group members may fall victim to a bad cascade, either informational or reputational. Deliberators may emphasize shared information at the expense of uniquely held information. Finally, group polarization may lead even rational people to unjustified extremism. By contrast, prediction markets often produce accurate results, because they create strong incentives for revelation of privately held knowledge and succeed in aggregating widely dispersed information. The success of prediction markets offers a set of lessons for increasing the likelihood that groups can obtain the information that their members have.

Many institutions, both public and private, make their decisions through deliberation. But why, exactly, is deliberation important or even desirable? A central answer must be that deliberation will result in wiser judgments and better outcomes. But does deliberation actually have this effect? The answer is by no means clear. Group members may impose pressures on one another, leading to a consensus on falsehood rather than truth. A group of like-minded people, with similar predilections, is particularly vulnerable to this problem. The idea of “groupthink,” coined and elaborated by Irving Janis, suggests the possibility that groups will tend toward uniformity and censorship, thus failing to combine information and enlarge the range of arguments.¹ Without structural protections, both private and public groups are likely to err, not in spite of deliberation but because of it.

My aim here is to compare deliberation with an intriguing social innovation—prediction markets—and to explore the advantages of the latter over the former in aggregating information. One of my goals is to see how the successes of prediction

^{*} Karl N. Llewellyn Distinguished Service Professor, Law School and Department of Political Science, University of Chicago. I am grateful to Christian List for valuable comments on a previous draft.

¹ Irving L. Janis, *Groupthink*, 2d ed. (Boston: Houghton Mifflin, 1982), 7–9.

markets might inform the practice of deliberation. To explain why deliberation often fails, I investigate two sets of influences on members of deliberating groups.² The first consists of informational influences, by which group members fail to disclose what they know out of deference to the information publicly announced by others. The second involves social pressures, which lead people to silence themselves in order not to face reputational sanctions, such as the disapproval of relevant others. As a result of these problems, groups often amplify rather than correct individual errors; emphasize shared information at the expense of unshared information; fall victim to cascade effects; and tend to end up in more extreme positions in line with the predeliberation tendencies of their members. In the United States, even federal judges are vulnerable to the relevant pressures, as both Republican and Democratic appointees show especially ideological voting when they are sitting with other judges appointed by presidents of the same political party.³

Because of these pressures, deliberative processes often fail to achieve their minimal goal of aggregating the information actually held by the deliberators. Indeed, such processes often fail to aggregate information even as they decrease variance, and increase confidence, among their members. A confident, cohesive, error-prone group is nothing to celebrate. On the contrary, it might be extremely dangerous, both to itself and to others.⁴

As we shall see, prediction markets often outperform deliberating groups, simply because they are so effective at pooling dispersed information among diverse people. Indeed, prediction markets realign private incentives in a way that makes them exceptionally well-designed to reduce the problems that infect deliberating groups. Such markets are worth investigating, in part because they provide an illuminating route by which to explore some characteristic defects in deliberative processes—and by which to

² I explore these mechanisms from different directions in Cass R. Sunstein, *Why Societies Need Dissent* (Cambridge, Mass.: Harvard University Press, 2003), and Cass R. Sunstein, *Infotopia: How Many Minds Aggregate Knowledge* (New York: Oxford University Press, 2006); I have borrowed from those accounts here.

³ Cass R. Sunstein, David Schkade, and Lisa Ellman, “Ideological Voting on Federal Courts of Appeals: A Preliminary Investigation,” *Virginia Law Review* 90 (2004): 301.

⁴ See the comparison of democratic and nondemocratic regimes in Dominic Johnson, *Overconfidence and War: The Havoc and Glory of Positive Illusions* (Cambridge, Mass.: Harvard University Press, 2004): 180–83.

obtain insights about how they might work better. In addition, such markets are worth investigating in their own right, if only because they promise to provide a supplement to deliberation that might well improve social decisions.

Deliberating Groups

If deliberating groups do well, we can imagine three principal reasons:

- *Groups are equivalent to their best members.* One or more group members will often know the right answer, and other members might well become convinced of this fact. For this reason, groups might perform toward or at the level of their best members. If some or many members suffer from ignorance or from a form of bias that leads to error, others might correct them. Deliberation might correct individual errors rather than propagate them, in a way that allows convergence on the judgment of the most accurate group member.
- *The whole is the sum of the parts: aggregating information.* Deliberation could aggregate existing information in a way that leads the group as a whole to know more than any individual member does. Suppose that the group contains no experts on the question at issue, but that relevant information is dispersed among members so that the group is potentially expert even if its members are not. Or suppose that the group contains a number of experts, but that each member is puzzled about how to solve a particular problem. Deliberation might elicit the relevant information and allow the group to make a sensible judgment. In this process, the whole is equal to the sum of the parts—and the sum of the parts is what is sought.
- *Improving on majority rule.* Suppose that in advance of deliberation, each group member is more than 50 percent likely to be right. The Condorcet Jury Theorem shows that the likelihood that the group's majority will be right expands to 100 percent as the size of the group increases. Perhaps deliberating groups will do better than would the majority of their individual members without deliberation—whatever the initial distribution of correct answers within those groups.
- *The whole goes beyond the sum of the parts: synergy.* The give and take of group discussion might sift information and perspectives in a way that leads the group to a good solution to a problem, one in which the whole is actually more than the sum of its parts. In such cases, deliberation is, at the very least, an ambitious form of information aggregation, one in which the exchange of views leads to a creative answer or solution.

To what extent do these mechanisms work in practice? Two points are entirely clear. First, deliberation usually reduces variance.⁵ After talking together, group members

⁵ Roger Brown, *Social Psychology*, 2d ed. (New York: Free Press, 1986): 206–7.

tend to come into accord with one another.⁶ Second, group members tend to become far more confident of their judgments after they speak with one another.⁷ A significant effect of group interactions is a greater sense that one's postdeliberation conclusion is correct—whether it actually is or not. Corroboration by others increases confidence in one's judgments.⁸ It follows that members of deliberating groups will usually converge on a position on which members have a great deal of confidence. This is not disturbing if that position is also likely to be correct—but if it is not, then many group members will end up sharing a view in which they firmly believe, but which turns out to be wrong (a most unfortunate and sometimes quite dangerous situation).

Unfortunately, there is no systematic evidence that deliberating groups will usually succeed in aggregating the information held by their members. With respect to questions with definite answers, deliberating groups tend to do about as well as or slightly better than their average members, but not as well as their best members.⁹ Hence, it is false to say that group members usually end up deferring to their internal specialists. Truth does not win out; the most that can be said is that under some conditions, the group will converge on the truth if the truth begins with “at least some initial support” within the group when the task has “a demonstrably correct answer.”¹⁰ Note here that when a group outperforms most of its individual members, it is generally because the issue is one on which a particular answer can be shown, to the satisfaction of all or most, to be right; and that even in that condition, the group might not do well if the demonstrably correct solution lacks significant support at the outset.

In general, simple majority schemes do fairly well at predicting group judgments for many decision tasks. It follows that if the majority is wrong, the group will be wrong

⁶ Ibid.

⁷ Chip Heath and Rich Gonzalez, “Interaction with Others Increases Decision Confidence but Not Decision Quality: Evidence against Information Collection Views of Interactive Decision Making,” *Organizational Behavior and Human Decision Processes* 61 (1995): 305.

⁸ See Robert Baron et al., “Social Corroboration and Opinion Extremity,” *Journal of Experimental Social Psychology* 32 (1996): 537.

⁹ Daniel Gigone and Reid Hastie, “Proper Analysis of the Accuracy of Group Judgments,” *Psychological Bulletin* 121 (1997): 149, 161; Reid Hastie, “Experimental Evidence of Group Accuracy,” in *Information Pooling and Group Decision Making*, ed. Bernard Grofman, Guillermo Owen, et al. (Greenwich, Conn.: JAI Press, 1986), 129.

¹⁰ Robert J. MacCoun, “Comparing Micro and Macro Rationality,” in *Judgments, Decisions, and Public Policy*, ed. Rajeev Gowda and Jeffrey Fox (Cambridge: Cambridge University Press, 2002), 116, 121.

as well.¹¹ With experts, the same general conclusion holds. Thus a “structured approach for combining independent forecasts is invariably more accurate” than “traditional group meetings,” which do “not use information efficiently.”¹²

Sources of Deliberative Failure

For two reasons, exposure to the views of others might lead people to silence themselves. The first involves the informational signals provided by the acts and views of other people. If most group members believe that X is true, there is reason to believe that X is in fact true, and that reason might outweigh the purely private reason a particular group member has to believe that X is false. If other group members share a particular belief, isolated or minority members might not speak out, deferring to the informational signal given by the statements of others. Not surprisingly, the strength of the signal will depend on the number and nature of the people who are giving it. People are particularly averse to being sole dissenters.¹³ If all but one person in a deliberating group has said that X is true, then the remaining member is likely to agree X is true, even to the point of ignoring the evidence of his own senses. And if the group contains one or more people who are well-known to be authorities, then other group members are likely to defer to them.

The second reason that group members might silence themselves involves social influences. Their silence might stem not from a belief that they are wrong, as in the case of informational pressure, but instead from the risk of social sanctions of various sorts. In the most extreme cases, those sanctions will take the form of criminal punishment or complete exclusion from the group. In less severe cases, those who defy the dominant position within the group will incur a form of disapproval that will lead them to be less trusted, liked, and respected in the future. Here, too, people are inevitably affected by the number and nature of those with the majority position. A large majority will impose more social pressure than a small one. If certain group members are leaders or authorities willing and able to impose social sanctions of various sorts, others will be unlikely to

¹¹ Ibid.

¹² J. Scott Armstrong, “Combining Forecasts,” in *Principles of Forecasting*, ed. J. Scott Armstrong (Boston: Kluwer Academic, 2001), 433.

¹³ See Sunstein, *Why Societies Need Dissent*.

defy them publicly.

Participation in deliberative processes, and the effects of informational and social influences, can be put into a more general framework. Suppose that group members are deliberating about some factual question; suppose, too, that each member has some information that bears on the answer to that question. Will members disclose what they know?

For each person, the answer may well depend on the individual benefits and the individual costs of disclosure. In many situations, and entirely apart from informational and social influences, the individual benefits of disclosure will be far less than the social benefits. In this sense, participants in deliberation often face a collective action problem, in which each person, following his rational self-interest, will tell the group less than it needs to know. At least, this is so if each member receives only a small portion of the benefits that come to the group from a good outcome—a plausible view about the situation facing many institutions, including, for example, labor unions, religious organizations, student and faculty groups, corporate boards, and government agencies

If the statements of others suggest that privately held information is wrong or unhelpful, then the private benefit of disclosure is reduced much more. In that event, the group member has reason to believe that disclosure will not improve the group's decision at all. Things are even worse if those who speak against the apparent consensus suffer reputational injury (or more). In that event, the private calculus is straightforward: Silence is golden.

Both informational pressure and social influences help explain the finding that in a deliberating group, those in a minority position often silence themselves or otherwise have disproportionately little weight. There is a more particular finding: Members of low-status groups—less-educated people, African-Americans, sometimes women—speak less and carry less influence within deliberating groups than their higher-status peers.¹⁴ Both informational influence and social pressures, likely to be especially strong for low-status members, contribute to this result. The unfortunate consequence can be a loss of information to the group as a whole, in a way that ensures that deliberating groups do far

¹⁴ Caryn Christenson and Ann Abbott, "Team Medical Decision Making," in *Decision Making in Health Care*, ed. Gretchen Chapman and Frank Sonnenberg (New York: Cambridge University Press, 2000), 267, 273–76.

less well than they would if only they could aggregate the information held by group members.

More generally, a comprehensive study has demonstrated that majority pressures can be powerful even for factual questions to which some people know the right answers.¹⁵ The study involved twelve hundred people, forming groups of six, five, and four members. Individuals were asked true-false questions involving art, poetry, public opinion, geography, economics, and politics. They were then asked to assemble into groups, which discussed the questions and produced answers. The majority played a substantial role in determining each group's answers. The truth played a role, too, but a lesser one. If a majority of individuals in the group gave the right answer, the group's decision moved toward the majority in 79 percent of the cases. If a majority of individuals in the group gave the wrong answer, the group's decision nonetheless moved toward the majority in 56 percent of the cases. Hence, the truth did have an influence—79 percent is higher than 56 percent—but the majority's judgment was the dominant one. And because the majority was influential even when wrong, the average group decision was right only slightly more often than the average individual decision (66 percent versus 62 percent). What is most important is that groups did not perform as well as they would have if they had properly aggregated the information that group members had.

Habermas vs. Hayek

Do these points amount to a challenge to deliberation as an ideal, or to deliberative conceptions of democracy? Many of those interested in deliberation have attempted to specify its preconditions in a way that is intended to ensure against predictable problems that infect real-world processes. Jürgen Habermas, for example, stresses norms and practices designed to allow victory by “the better argument”:

Rational discourse is supposed to be public and inclusive, to grant equal communication rights for participants, to require sincerity and to diffuse any kind of force other than the forceless force of the better argument. This communicative structure is expected to create a deliberative space for the mobilization of the best available contributions for the most relevant topics.¹⁶

¹⁵Robert L. Thorndike, “The Effect of Discussion upon the Correctness of Group Decisions: When the Factor of Majority Influence Is Allowed For,” *Journal of Social Psychology* 9 (1938): 343.

¹⁶See Jürgen Habermas, *Between Facts and Norms: An Author's Reflections*, 76 DENV. U. L. REV. 937, 940 (1999)..

In Habermas's "ideal speech situation," all participants attempt to seek the truth; they do not behave strategically or attempt to decide; they accept a norm of equality.¹⁷ Other advocates of deliberative democracy have spoken similarly about what appropriate deliberation entails.¹⁸ On this view, deliberation, properly understood, does not simply involve the exchange of words and opinions. It imposes its own requirements and preconditions. Indeed, deliberation has its own internal morality, one that operates as a corrective to some of the effects of deliberative processes in the real world.

Unfortunately, preconditions of the sort identified by Habermas will cure few of the problems that I shall be outlining here. Those preconditions will do little to affect the key failures on the part of deliberating groups. Each of the failures is likely to arise even if discourse is public and inclusive, even if participants are sincere, and even if everyone has equal communication rights. We might therefore take the argument here as a Hayekian challenge to Habermas—a challenge that stresses (with Friedrich Hayek) the diffusion of information in society and the difficulty of aggregating that information through deliberation (as opposed to the price signal, which Hayek championed).¹⁹

Consider four sets of deliberative failures.

Deliberative Failure 1: Amplification of Cognitive Errors

It is well known that individuals do not always process information well. They use heuristics that lead them to predictable errors; they are also subject to identifiable biases, which also produce errors.²⁰ For example, most people follow the representativeness heuristic, in accordance with which judgments of probability are influenced by assessments of resemblance (the extent to which A "looks like" B).²¹ The representativeness heuristic helps explain what Paul Rozin and Carol Nemeroff have

¹⁷ See JÜRGEN HABERMAS, *What is Universal Pragmatics?*, in COMMUNICATION AND THE EVOLUTION OF SOCIETY 1, 2–4, 32 (Thomas McCarthy trans., 1979) (discussing preconditions for communication).

¹⁸ See Amy GUTMANN & DENNIS THOMPSON, *Democracy and Disagreements* 7–8 (1997) (outlining foundations of authors' vision of deliberative democracy).

¹⁹ See generally F.A. Hayek, *The Use of Knowledge in Society*, 35 AM. ECON. REV. 519 (1945) (discussing dispersal of knowledge and its aggregation through markets).

²⁰ For an overview, see Thomas Gilovich, Dale Griffin, and Daniel Kahneman, *Heuristics and Biases: The Psychology of Intuitive Judgment* (New York: Cambridge University Press, 2002).

²¹ Amos Tversky and Daniel Kahneman, "Judgment under Uncertainty: Heuristics and Biases," in *Judgment under Uncertainty: Heuristics and Biases*, ed. Daniel Kahneman, Paul Slovic, and Amos Tversky (Cambridge: Cambridge University Press), 1982, 3.

called “sympathetic magical thinking,” including the beliefs that some objects have contagious properties, and that causes resemble their effects.²² The representativeness heuristic often works well, but it can also lead to severe blunders.

People often err because they use the availability heuristic to answer difficult questions about probability. When people use the availability heuristic, they answer a question of probability by asking whether examples come readily to mind.²³ Consider, for example, the question whether we should fear a hurricane, a nuclear power accident, or a terrorist attack. If it is easy to think of a case in which one of these hazards created serious harm, the assessment of probability will be greatly affected. Of course, use of the availability heuristic is not irrational, but it, too, can produce both excessive and insufficient fear.

For purposes of assessing deliberation, a central question is whether groups avoid the errors of the individuals who comprise them. There is no clear evidence that they do, and often they do not—a vivid illustration of the principle, “garbage in, garbage out,” in a way that mocks the aspiration to collective correction of individual blunders. In fact, individual errors are not merely replicated but actually amplified in group decisions—a process of “some garbage in, much garbage out.”

Consider some key findings. If individual jurors are biased because of pretrial publicity that misleadingly implicates the defendant, or even because of the defendant’s unappealing physical appearance, juries are likely to amplify rather than correct those biases.²⁴ Groups have been found to amplify, rather than to attenuate, reliance on the representativeness heuristic;²⁵ to reflect even larger framing effects than individuals;²⁶ to show more overconfidence than group members;²⁷ to be more affected

²² Paul Rozin and Carol Nemeroff, “Sympathetic Magical Thinking: The Contagion and Similarity ‘Heuristics,’” in Gilovich, Griffin, and Kahneman, *Heuristics and Biases* p. 201

²³ Tversky and Kahneman, “Judgment under Uncertainty,” 3.

²⁴ MacCoun, “Comparing Micro and Macro Rationality,” 116, 121.

²⁵ Mark F. Stasson et al., “Group Consensus Approaches in Cognitive Bias Tasks,” *Japanese Psychological Research* 30 (1988): 68.

²⁶ See Norbert L. Kerr et al., “Bias in Judgment: Comparing Individuals and Groups,” *Psychology Review* 103 (1996): 687, 689, 691–93.

²⁷ Janet A. Snizek and Rebecca A. Henry, “Accuracy and Confidence in Group Judgment,” *Organizational Behavior and Human Decision Processes* 43 (1989): 1. This finding very much bears on excessive risk-taking, including in the context of making war. See Dominic Johnson, *Overconfidence and War*, 180–83.

by the biasing effect of spurious arguments from lawyers;²⁸ to be more susceptible to the “sunk cost fallacy”;²⁹ and to be more subject to choice-rank preference reversals.³⁰ In an especially revealing finding, groups have been found to make more, rather than fewer, conjunction errors than individuals when individual error rates are high—though fewer when individual error rates are low.³¹ In addition, groups demonstrate essentially the same level of reliance on the availability heuristic, even when use of that heuristic leads to clear errors.³²

Deliberative Failure 2: Hidden Profiles and Common Knowledge

Suppose that group members have a great deal of information—enough to produce the unambiguously right outcome if that information is properly aggregated. Even if this is so, an obvious problem is that groups will not perform well if they emphasize shared information and slight information that is held by one or a few members. Unfortunately, countless studies demonstrate that this regrettable result is highly likely.³³ “Hidden profiles” is the term for accurate understandings that groups could but do not obtain. Hidden profiles are, in turn, a product of the *common-knowledge effect*, through which information held by all group members has more influence on group judgments than information held by only a few members.³⁴ The most obvious explanation of the effect is the simple fact that as a statistical matter, common knowledge is more likely to be communicated to the group; but social influences play a role as well.

Hidden Profiles. Consider a study of serious errors within working groups, both

²⁸ Edward L. Schumann and W. C. Thompson, “Effects of Attorney’s ^{**}{ {AU:OK, or “Attorneys”} } Arguments on Jurors’ Use of Statistical Evidence” (unpublished manuscript, 1989).

²⁹ Glen Whyte, “Escalating Commitment in Individual and Group Decision Making,” *Organizational Behavior and Human Decision Processes* 54 (1993): 430.

³⁰ James W. Gentry and John C. Mowen, “Investigation of the Preference Reversal Phenomenon in a New Product Introduction Task,” *Journal of Applied Psychology* 65 (1980): 715; Julie R. Irwin and James H. Davis, “Choice/Matching Preference Reversals in Groups,” *Organizational Behavior and Human Decision Processes* 64 (1995): 325.

³¹ Whyte, “Escalating Commitment,” 430.

³² Stasson et al., “Group Consensus Processes,” ^{**}{ {AU:Please check—title in earlier note is “Group Consensus Approaches”} } 68.

³³ Garold Stasser and William Titus, “Hidden Profiles: A Brief History,” *Psychological Inquiry* 14 (2003): 304.

³⁴ Daniel Gigone and Reid Hastie, “The Common Knowledge Effect: Information Sharing and Group Judgments,” *Journal of Personality and Social Psychology* 65 (1993): 959.

face-to-face and online.³⁵ The purpose of the study was to see how groups might collaborate to make personnel decisions. Resumes for three candidates applying for a marketing manager position were placed before group members. The attributes of the candidates were rigged by the experimenters so that one applicant was clearly the best for the job described. Packets of information were given to subjects, each containing a subset of information from the resumes, so that each group member had only part of the relevant information. The groups consisted of three people, some operating face-to-face, some operating online. Almost none of the deliberating groups made what was conspicuously the right choice. The reason is simple: They failed to share information in a way that would permit the group to make that choice. Members tended to share positive information about the winning candidate and negative information about the losers. They suppressed negative information about the winner and positive information about the losers. Hence, their statements served to “reinforce the march toward group consensus rather than add complications and fuel debate.”³⁶

Or consider a simulation of political elections, in which information was parceled out to individual members about three candidates for political office, and in which properly pooled information could have led to what was clearly the best choice, candidate A.³⁷ In the first condition, each member of the four-person groups was given most of the relevant information (66 percent of the information about each candidate). In that condition, 67 percent of group members favored candidate A before discussion and 85 percent after discussion.³⁸ This is a clear example of appropriate aggregation of information. Groups significantly outperformed individuals, apparently because of the exchange of information and reasons. Here, then, is a clear illustration of the possibility that groups can aggregate what members know in a way that produces sensible outcomes.

In the second condition, by contrast, the information that favored candidate A was parceled out to various members of the group so that only 33 percent of information

³⁵ See Ross Hightower and Lutfus Sayeed, “The Impact of Computer-Mediated Communication Systems on Biased Group Discussion,” *Computers in Human Behavior* 11 (1995): 33.

³⁶ Patricia Wallace, *The Psychology of the Internet* (Cambridge: Cambridge University Press, 1999), 82.

³⁷ See Garold Stasser and William Titus, “Pooling of Unshared Information in Group Decision Making: Biased Information Sampling during Discussion,” *Journal of Personality and Social Psychology* 48 (1985): 1467.

³⁸ *Ibid.*, 1473; see also Stasser and Titus, “Hidden Profiles,” 304.

about each candidate was shared. As the condition was designed, the shared information favored two unambiguously inferior candidates, B and C; but if the unshared information emerged through discussion, and were taken seriously, candidate A would be chosen. In that condition, less than 25 percent of group members favored candidate A before discussion, a natural product of the initial distribution of information. But (and this is the key result) that number actually *fell* after discussion, simply because the shared information had disproportionate influence on group members.³⁹ In other words, groups did worse, not better, than individuals when the key information was distributed selectively. In those conditions, the commonly held information was far more influential than the distributed information, to the detriment of the group's ultimate decision.

From this and many similar studies, the general conclusion is that when "the balance of unshared information opposes the initial most popular position . . . the unshared information will tend to be omitted from discussion and, therefore, will have little effect on members' preferences during group discussion."⁴⁰ It follows that "group decisions and postgroup preferences reflect the initial preferences of group members even when the exchange of unshared information should have resulted in substantial shifts in opinion."⁴¹ Nor does discussion increase the recall of unshared information. On the contrary, its major effect is to increase recall of the attributes of the initially most popular candidate.⁴² The most disturbing conclusion is that when key information is unshared, groups are "more likely to endorse an inferior option after discussion than [are] their individual members before discussion."⁴³

The Common-Knowledge Effect. These results are best understood as a consequence of the common-knowledge effect, by which information held by all group members has far more influence on group judgments than information held by one member or a few.⁴⁴ More precisely, the "influence of a particular item of information is directly and positively related to the number of group members who have knowledge of

³⁹ Stasser and Titus, "Pooling of Unshared Information," 1473.

⁴⁰ *Ibid.*, 1476.

⁴¹ *Ibid.*

⁴² *Ibid.*

⁴³ Stasser and Titus, "Hidden Profiles," 305.

⁴⁴ See Daniel Gigone and Reid Hastie, "The Common Knowledge Effect: Information Sharing and Group Judgments," *Journal of Personality and Social Psychology* 65 (1993): 959.

that item before the group discussion and judgment.”⁴⁵ Under conditions of unshared information, group judgments have been found to be “not any more accurate than the average of the individual judgments, even though”—and this is the central point—the groups were “in possession of more information than were any of the individuals.”⁴⁶

As might be expected, the group’s focus on shared information increases with the size of the group.⁴⁷ In a study by Stasser and colleagues designed to test judgments about candidates for office, involving both three-person and six-person groups, all discussions focused far more on shared than on unshared information—but the effect was significantly greater for six-person groups. Most remarkably, the researchers write, “it was almost as likely for a shared item to be mentioned twice as it was for an unshared item to be mentioned at all.”⁴⁸ And despite the failures of their deliberations, group members were significantly more confident in their judgments after discussion.⁴⁹

Deliberative Failure 3: Cascades

A cascade is a process by which people influence one another, so much so that participants ignore their private knowledge and rely instead on the publicly stated judgments of others. There are two kinds of cascades: informational and reputational. In informational cascades, people silence themselves out of deference to the information conveyed by others. In reputational cascades, they silence themselves so as to avoid the opprobrium of others.

Informational Cascades. Hidden profiles are closely related to informational cascades, which greatly impair group judgments. Cascades need not involve deliberation, but deliberative processes often involve cascades. As in the case of hidden profiles, the central point is that those involved in a cascade do not reveal what they know. As a result, the group does not obtain important information.

To see how informational cascades work, imagine a deliberating group that is

⁴⁵ Ibid., 960.

⁴⁶ Ibid., 973.

⁴⁷ See Garold Stasser et al., “Information Sampling in Structured and Unstructured Discussions of Three and Six-Person Groups,” *Journal of Personality and Social Psychology* 57 (1989): 67.

⁴⁸ Ibid., 78.

⁴⁹ Ibid., 72.

deciding whether to authorize some new venture.⁵⁰ Let us also assume that the jurors are announcing their views in sequence, in a temporal queue, and that each member knows his place in that queue. Every member has some private information about what should be done. But each also attends, reasonably enough, to the judgments of others. Mr. Andrews is the first to speak. He suggests that the venture should be authorized. Ms. Barnes now knows Andrews's judgment; it is clear that she, too, should vote in favor of the venture if she agrees independently with Andrews. But if her independent judgment is otherwise, she would—if she trusts Andrews no more and no less than she trusts herself—be indifferent about what to do and might simply flip a coin.

Now turn to a third juror, Mr. Carlton. Suppose that both Andrews and Barnes have argued in favor of the venture but that Carlton's own information, though inconclusive, suggests that the venture is a terrible idea. In that event, Carlton might well ignore what he knows and follow Andrews and Barnes. It is likely in these circumstances that both Andrews and Barnes had reasons for their conclusion, and unless Carlton thinks that his own information is better than theirs, he should follow their lead. If he does, Carlton is in a cascade. Now suppose that Carlton is acting in response to what Andrews and Barnes did, not on the basis of his own information, and that subsequent members know what Andrews, Barnes, and Carlton did. On reasonable assumptions, they will do exactly what Carlton did: favor the venture regardless of their private information (which, we are supposing, is relevant but inconclusive). This will happen even if Andrews initially blundered.⁵¹

If this is what is happening, there is a serious social problem: Those who are in the cascade do not disclose the information that they privately hold. In the example just given, decisions will not reflect the overall knowledge, or the aggregate knowledge, of those in the group—even if the information held by individual members, if actually revealed and aggregated, would produce a quite different result. The reason is that people are following the lead of those who came before. Subsequent speakers might fail to rely on, and fail to reveal, private information that actually exceeds the information

⁵⁰ I draw here on David Hirschleifer, "The Blind Leading the Blind," in *The New Economics of Human Behavior*, ed. Marianno Tommasi and Kathryn Ierulli (Cambridge: Cambridge University Press, 1995), 188, 193–94.

⁵¹ *Ibid.*, 195.

collectively held by those who started the cascade.

Cascades often occur in the real world within deliberating groups or elsewhere;⁵² they are easy to create in the laboratory. The simplest experiment asked subjects to guess whether the experiment was using urn A, which contained two red balls and one white, or urn B, which contained two white balls and one red.⁵³ Subjects could earn \$2.00 for a correct decision, and hence an economic incentive favored correct individual decisions (a point to which I will return). In each period, the contents of the chosen urn were emptied into a container. A randomly selected subject was asked to make one (and only one) private draw of a ball in each round. The subject recorded the color of that draw on an answer sheet and his own decision about which urn was involved. The subject did not announce his draw to the group, but he did announce his own decision to everyone. Then the urn was passed to the next subject for his own private draw, which again was not disclosed, and his own decision about the urn, which again was disclosed. This process continued until all subjects had made draws and decisions. At that time, the experimenter announced the actual urn used. If the subject had picked the urn only on the basis of his private information, he would have been right 66.7 percent of the time. The point of the experiment was to see whether people will decide to ignore their own draw in the face of conflicting announcements by predecessors—and to explore whether such decisions will lead to cascades and errors.

In the experiment, cascades often developed and often produced errors. After a number of individual judgments were revealed, people sometimes announced decisions that were inconsistent with their private draws, but that fit with the majority of previous announcements.⁵⁴ More than 77 percent of “rounds” resulted in cascades, and 15 percent of private announcements did not reveal a “private signal,” that is, the information provided by people’s own draws. Consider cases in which one person’s draw (say, red) contradicted the announcement of his predecessor (say, urn B). In such cases, the second announcement nonetheless matched the first about 11 percent of the time—far less than a

⁵² See *ibid*; also see Sunstein, *Why Societies Need Dissent*.

⁵³ See Lisa Anderson and Charles Holt, “Information Cascades in the Laboratory,” *American Economic Review* 87 (1997): 847.

⁵⁴ See Angela Hung and Charles Plott, “Information Cascades: Replication and an Extension to Majority Rule and Conformity-Rewarding Institutions,” *American Economic Review* 91 (2001): 1508, 1515.

majority, but enough to ensure cascades. And when one person's draw contradicted the announcement of two or more predecessors, the second announcement was likely to follow those who went before. Of note, the majority of decisions were rationally based on the available information⁵⁵—but erroneous cascades nonetheless developed. Table 4-1 shows an example of a cascade that produced an inaccurate outcome (the urn used was B):⁵⁶

Table 4-1. An Informational Cascade

	1	2	3	4	5	6
Private Draw	a	A	a	b	b	b
Decision	A	A	A	A	A	A

Source: Willinger and Ziegelmeyer, "Are More Informed Agents," 291.

What is noteworthy here, of course, is that the total amount of private information—four whites and two reds—justified the correct judgment, which was in favor of urn B. But the existence of two early signals, producing rational but incorrect judgments, led everyone else to fall in line. "Initial misrepresentative signals start a chain of incorrect decisions that is not broken by more representative signals received later."⁵⁷ This result maps directly onto real-world decisions by deliberating groups, in which people fail to disclose what they know, to the detriment of the group as a whole.

Reputational Cascades. In a reputational cascade, people think they know what is right, or what is likely to be right, but they nonetheless go along with the crowd in order to maintain the good opinion of others. Suppose Albert suggests that global warming is a serious problem and that Barbara concurs with Albert, not because she actually thinks that Albert is right, but because she does not wish to seem, to Albert, ignorant or indifferent to environmental protection. If Albert and Barbara seem to agree that global warming is a serious problem, Cynthia not only might not contradict them publicly, but also might even appear to share their judgment, not because she believes that judgment to be correct, but because she does not want to face their hostility or lose

⁵⁵ Thus, 72 percent of subjects followed Bayes's rule in Lisa Anderson and Charles Holt, "Information Cascades in the Laboratory," *American Economic Review* 87 (1997): 847, and 64 percent in Marc Willinger and Anthony Ziegelmeyer, "Are More Informed Agents Able to Shatter Information Cascades in the Lab?" in *The Economics of Networks: Interaction and Behaviours*, ed. Patrick Cohendet et al. (New York: Springer, 1998), 291, 304.

⁵⁶ See Willinger and Ziegelmeyer, "Are More Informed Agents," 291.

⁵⁷ See Anderson and Holt, "Information Cascades in the Laboratory," 847.

their good opinion.

It should be easy to see how this process might generate a cascade. Once Albert, Barbara, and Cynthia offer a united front on the issue, their friend David might be most reluctant to contradict them, even if he thinks they are wrong. In the actual world of group decisions, people are, of course, uncertain whether publicly expressed statements are a product of independent information, participation in an informational cascade, or reputational pressure. Much of the time, listeners and observers undoubtedly overstate the extent to which the actions of others are based on independent information.

The possibility of reputational cascades is demonstrated by an ingenious variation on the urn experiment mentioned above.⁵⁸ In this experiment, people were paid \$0.25 for a correct decision, but \$0.75 for a decision that matched the decision of the majority of the group. There were punishments for incorrect and nonconforming answers as well. If people made an incorrect decision, they lost \$0.25; if their decision failed to match the group's decision, they lost \$0.75.

In this experiment, cascades appeared almost all of the time. No fewer than 96.7 percent of rounds resulted in cascades, and 35.3 percent of people's announcements did not match their private signal, that is, the signal given by their own draw. And when the draw of a subsequent person contradicted the announcement of the predecessor, 72.2 percent of people matched the first announcement. Consider, as a dramatic illustration, table 4-2, which shows this period of the experiment (the actual urn was B):⁵⁹

Table 4-2. Conformity and Cascades

	1	2	3	4	5	6	7	8	9	10
Private Draw	a	b	b	b	a	b	b	b	a	B
Decision	A	A	A	A	A	A	A	A	A	A

Source: Hung and Plott, "Information Cascades."

This experiment shows that especially unfortunate results should be expected if people are rewarded not only or not mostly for being correct, but also or mostly for doing what other people do. The problem is that people are not revealing the information they actually have.

⁵⁸ See Hung and Plott, "Information Cascades," 1515–17.

⁵⁹ Ibid., 1516.

Deliberative Failure 4: Group Polarization

There are clear links among hidden profiles, social cascades, and the well-established phenomenon of group polarization, by which *members of a deliberating group end up adopting a more extreme version of the position toward which they tended before deliberation began*.⁶⁰ The problem is especially severe for groups of like-minded people, who typically end up in more extreme positions as a result of deliberation. Group polarization is the typical pattern with deliberating groups, and it has been found in hundreds of studies involving more than a dozen countries, including the United States, France, Afghanistan, and Germany.⁶¹ For example, those who disapprove of the United States and are suspicious of its intentions will increase their disapproval and suspicion if they exchange points of view. Indeed, there is specific evidence of the latter phenomenon among citizens of France.⁶²

Group polarization occurs for matters of fact as well as issues of value, though it is easier to demonstrate the latter. If the question is whether a terrorist attack will occur in the United States in the next year, group polarization will not be easy to test, simply because the answer is either yes or no, and it is not simple to demonstrate greater extremism in binary choices. But suppose that people are asked, on a bounded scale of zero to eight, how likely it is that a terrorist attack will occur in the United States in the next year, with zero indicating “zero probability,” eight indicating “absolutely certain,” seven indicating “overwhelmingly likely,” six “more probable than not,” and five “fifty-fifty.” In that event, the answers from a deliberating group will tend to reveal group polarization, as people move toward more extreme points on the scale depending on their initial median point. If the predeliberation median is five, the group judgment will usually be six; if the predeliberation median is three, the group judgment will usually be two.⁶³ Recall here that federal judges are highly susceptible to group polarization, as both Democratic and Republican appointees show far more ideological voting patterns when sitting with other judges appointed by a president of the same political party.⁶⁴ Juries

⁶⁰ Roger Brown, *Social Psychology: The Second Edition* (New York, N.Y.: Free Press, 1986), 206–7.

⁶¹ *Ibid.*, 204.

⁶² *Ibid.*, 224.

⁶³ *Ibid.*

⁶⁴ Sunstein, Schkade, and Ellman, “Ideological Voting,” 301.

polarize as well.⁶⁵

Why does group polarization occur? There are three reasons.⁶⁶ The first and most important involves the now-familiar idea of informational influence, but in a distinctive form. People respond to the arguments made by other people—and the “argument pool” in any group with some predisposition in one direction will inevitably be skewed toward that predisposition. As a statistical matter, the arguments favoring the initial position will be more numerous than those pointing in the other direction. Individuals will have heard of some, but not all, of the arguments that emerge from group deliberation. As a result of the relevant arguments, deliberation will lead people toward a more extreme point in line with what group members initially believed.

The second explanation involves social influences. People want to be perceived favorably by other group members. Sometimes people’s publicly stated views are, to a greater or lesser extent, a function of how they want to present themselves. Once they hear what others believe, some will adjust their positions at least slightly in the direction of the dominant position in order to hold onto their preserved self-presentation. They shift accordingly.⁶⁷

The third explanation stresses that people with extreme views tend to have more confidence that they are right, and that, as people gain confidence, they become more extreme in their beliefs.⁶⁸ In a wide variety of experimental contexts, people’s opinions have been shown to become more extreme simply because their views have been corroborated and because they have been more confident after learning of the shared views of others.⁶⁹

Note that if it is understood in these terms, group polarization may well reflect rational behavior at the individual level.⁷⁰ Suppose that each group member privately assigns a significant probability, say of 0.6, to the truth of some hypothesis (say, that

⁶⁵ See David Schkade et al., “Deliberating about Dollars: The Severity Shift,” *Columbia Law Review* 100 (2000): 101.

⁶⁶ Brown, *Social Psychology*, 200–45.

⁶⁷ Ibid. It has similarly been suggested that majorities are especially potent because people do not want to incur the wrath, or lose the favor, of large numbers of others, and that when minorities have influence, it is because they produce genuine attitudinal change. See Baron et al., “Social Corroboration,” 82.

⁶⁸ Baron et al., “Social Corroboration,” 537.

⁶⁹ Ibid.

⁷⁰ I am grateful to Christian List for pressing this point; he should not be held responsible for my restatement of it here.

North Korea will have nuclear weapons within the next year). Suppose, further, that group discussion leads each group member to hear evidence largely supportive of the hypothesis, leading to a judgment in favor of a higher probability, say of 0.7. Rational updating may be entirely responsible for the shift. Now suppose all group members report after deliberation that they have each independently arrived at a probability of 0.7 for the truth of the hypothesis, based on the evidence they have received (and starting with a prior of 0.6). What this means is that there have been a total number of n independent items of evidence in support of the hypothesis (one such item for each of n group members), each of which has been sufficiently strong to support a Bayesian update from 0.6 to 0.7. The existence of these n independent items of evidence should then lead the group as a whole—and each group member post-deliberation—to assign a still higher posterior probability to the hypothesis, i.e., a probability well above 0.7.

Whether rational updating of this kind will produce accurate or inaccurate judgments depends on the antecedently held information within the group. Suppose that for people who are fully informed, the probability that the relevant hypothesis is true is in fact 0.3. If the group starts from a significantly inflated probability estimate, group polarization will lead them to make severe errors. Nothing in the phenomenon of group polarization demonstrates that deliberation will lead to blunders. But if the median predeliberation view is wrong, groups are likely to do worse than their individual members.

It should be clear that the four sources of deliberative failure can create serious problems for deliberating groups. What might be done in response? I have mentioned Hayek's suggestion that the price mechanism is an excellent way to aggregate dispersed information. Might the price system be enlisted as a supplement to, or even a replacement for, social deliberation?

Prediction Markets

Deliberation is one way to aggregate privately held information, but there are many other possibilities. An obvious alternative is to rely on the price signal, which has a similar aggregative function. As Hayek emphasized, the price mechanism is a kind of “marvel,” because it combines widely dispersed information held by diverse people. And

if an emphasis is placed on the information-aggregating properties of markets, it would seem plain that, to improve on the answer produced by deliberating groups, we might consider an increasingly popular possibility: *Create a market*.⁷¹ Prediction markets, a recent innovation, have proved remarkably successful at forecasting future events; they seem to do far better, in many domains, than deliberating groups. Such markets are worth sustained attention, in part because they offer important lessons about how to make deliberation go better or worse, and in part because they provide a useful model for many private and public organizations.

Potential and Promise. A central advantage of prediction markets is that they impose the right incentives for diverse people to disclose the information they separately hold. Recall that in a deliberating group, members often have little incentive to say what they know. By speaking out, they provide benefits to others while possibly facing high private costs. Prediction markets realign incentives in a way that is precisely designed to overcome these problems. Because investments in such markets are generally not disclosed to the public, investors need not fear reputational sanctions if, for example, they have predicted that a company's sales will be low or that a certain candidate will be elected president. And because people stand to gain or lose from their investments, they have a strong incentive to use (and in that sense to disclose) whatever private information they hold; they can capture, rather than give to others, the benefits of disclosure. The use of private information will be reflected in the price signal. In these crucial ways, the problems that infect deliberating groups are largely eliminated in prediction markets.

Prediction markets also impose strong incentives for traders to ferret out accurate information. Traders do not trade blindly, and they are entirely able to stop trading, for a moment or more, in order to retrieve better information that will give them an advantage. In many deliberating groups, by contrast, participants cannot leave; they must continue deliberating, and the necessary information is, at best, dispersed and locked within individual participants. Well-functioning systems of deliberation encourage group members to act dynamically to acquire further information, just as markets do.

⁷¹ For valuable overviews, see Justin Wolfers and Eric Zitzewitz, "Prediction Markets," *Journal of Economic Perspectives* 18 (2004): 107; Michael Abramowicz, "Information Markets, Administrative Decisionmaking, and Predictive Cost-Benefit Analysis," *University of Chicago Law Review* 71 (2004): 933; Saul Levmore, "Simply Efficient Markets and the Role of Regulation," *Journal of Corporation Law* 28 (2003): 589.

Of course, investors, like everyone else, are subject to the informational pressure imposed by the views of others. But a market creates strong incentives for revelation of whatever information people actually hold. And indeed, prediction markets have been found not to amplify individual errors but to eliminate them; the prices that result from trading prove reliable even if many individual traders err. In recent years, prediction markets have done more than to provide valuable information. In countless domains, their forecasts have proved extremely accurate.⁷² The most dramatic finding is that prices generally operate as probabilities. When prices suggest that events are likely to occur with 90 percent probability, they occur 90 percent of the time; when the price suggest a probability of 80 percent, the events happen 80 percent of the time; and so forth.

Since 1988, the University of Iowa has run the Iowa Electronic Markets (IEM), which allow people to bet on the outcome of presidential elections. Before the 2004 elections, they did far better than professional polling organizations, outperforming polls 451 out of 596 times.⁷³ In the week before the four elections from 1988 to 2000, the predictions in the Iowa market showed an average absolute error of just 1.5 percentage points—a significant improvement over the 2.1 percentage point error in the final Gallup polls. In 2004, the Iowa market did even better. On midnight of November 1, it showed Bush with 50.45% of the vote and Kerry with 49.55%—very close to the final numbers of 51.56% for Bush and 48.44% for Kerry.

Prediction markets, aggregating diverse views, are flourishing in numerous domains. Consider the Hollywood Stock Exchange, in which people predict (among other things) Oscar nominees and winners as well as opening weekend box office figures. The level of accuracy has been extremely impressive, especially in view of the fact that the traders use virtual rather than real money. Among the most impressive achievements of the Hollywoods Stock exchange to date is its uncanny accuracy in predicting Oscar winners, with correct judgments in fifteen of the sixteen the categories for which trading

⁷² For early overviews, see David M. Pennock et al., *The Real Power of Artificial Markets*, 291 *Science* 987 (2001); David M. Pennock et al., *The Power of Play: Efficiency and Forecast Accuracy in Web Market Games*, NEC Research Institute Technical Report 1000-168 (2001). A wealth of valuable information can be found at <http://www.chrisfmasse.com/>

⁷³ See Robert W. Hahn & Paul C. Tetlock, *Harnessing the Power of Information: A New Approach to Economic Development* 4 (AEI-Brookings Joint Ctr. For Regulatory Studies, Working Paper No. 04-21, 2004), available at <http://www.aei-brookings.org/publications/abstract.php?pid=846>.

was allowed in the last two years. The markets for the demand for gas outperform the experts on the demand for gas.⁷⁴ Many people believe that “you can’t predict the weather,” but the National Weather Service does quite well, and Orange Juice futures do even better.⁷⁵ A large prediction market focuses on the likelihood that economic data released later in the week will show specific values;⁷⁶ the market has performed even better than the consensus forecasts of a survey of about fifty professional forecasters.

Many companies are now using prediction markets to aggregate diverse views. Hewlett Packard (HP) and the California Institute of Technology initiated a project to study prediction markets as an information aggregation mechanism involving product sales.⁷⁷ In no fewer than six of the eight markets for which official forecasts were available, the market prediction was significantly closer to the actual outcome than the official forecast. For its part, Google has created a large set of prediction markets to help to forecast its own development.⁷⁸ The relevant markets predict launch dates for products, new office openings, and a range of other outcomes of importance to the company. The outcomes have been exceedingly accurate; prices have actually represented probabilities. Dispersed knowledge within the company has been accurately aggregated in this way. Many other companies, including Ely Lilly and Microsoft, have used prediction markets as well to supplement deliberation about future courses of action.

To be sure, prediction markets themselves involve a measure of deliberation. Many individual investors are likely to have deliberated with others before they invest. In some such markets, investors undoubtedly act as “teams,” pooling resources after deliberating together about what to do. The point is that decisions ultimately come not from asking group members to come up with a mutually agreeable conclusion, but by reference to the price signal, which will have aggregated a great deal of diverse information. It is for this reason that prediction markets outperform deliberative

⁷⁴ See Robin Hanson, *Designing Real Terrorism Futures*, *supra* note, at 2.

⁷⁵ Richard Roll, *Orange Juice and Weather*, 74 *Am. Econ. Rev.* 861, 871 (1984).

⁷⁶ See Wolfers & Zitzewitz, *supra* note 207, at 113–14.

⁷⁷ See Kay-Yut Chen & Charles R. Plott, *Information Aggregation Mechanisms: Concept, Design, and Implementation for a Sales Forecasting Problem 3* (Div. of the Humanities & Soc. Sci., Cal. Inst. of Tech., Social Science Working Paper No. 113, March 2002) (describing variation of this model employed by Hewlett-Packard), available at <http://www.hss.caltech.edu/SSPapers/wp1131.pdf>.

⁷⁸ See Putting Crowd Wisdom to Work, available at <http://googleblog.blogspot.com/2005/09/putting-crowd-wisdom-to-work.html>

processes.

Strategy, Manipulation, and Limitations. It is natural to wonder about whether and when prediction markets might fail. Some clues are provided by two conspicuous failures. Such markets found it more probable than not that Special Prosecutor Patrick Fitzgerald would indict White House adviser Karl Rove in 2005, and they found it exceedingly improbable that President George W. Bush would appoint John Roberts to the United States Supreme Court. The best explanation is that there was not a great deal of dispersed information about the particular decisions of Special Prosecutor Fitzgerald or President Bush. To be sure, investors knew that Fitzgerald would not indict Bush himself and that President Bush would not appoint Fitzgerald (or Tony Blair, Saddam Hussein, or John Kerry) to the Supreme Court; but they lacked the kind of information that would permit successful judgments about the probability that a particular movie would win the Oscars, or that a particular product would do well in the market, or that a particular candidate would be elected in a contested race.

There is an additional problem. Suppose that investors know that their “bets” might have a significant impact on the hypothesis that the market is supposed to predict. Investors might believe, for example, that the predictions of the Iowa Electronic Markets will affect the outcomes of political campaigns, by making certain candidates look promising or instead doomed. Such investors might pour immense sums of money into bets on their preferred candidates, and in the process hope to create a self-fulfilling prophecy. This danger seems nonexistent for genuinely exogenous events, such as a natural disaster or an unusual weather pattern. But for economic events, political campaigns, product success, and terrorism, the risk of manipulation cannot be ruled out of bounds.

Existing evidence does suggest that the risk may be more hypothetical than real. Several efforts to manipulate election markets have been made, and they have not succeeded: In a short time, canny investors see that prices are inflated or deflated, and the price rapidly returns to normal. More experience is required to know whether manipulation will work in other contexts.

Feasibility, Markets, and Deliberation Once More. I have suggested that prediction markets face a pervasive problem of feasibility. A deliberating jury, for

example, could not enlist such markets to decide on questions of guilt or innocence. Among other things, there is no objective way to test whether the jury, or individual jurors, ended up with the right answer (and if there were, the jury might well be dispensable). More generally, it is not easy to see how prediction markets could be used on normative questions. At most, such markets could be used on the factual questions that are sometimes part of such questions.

There is another problem. When the relevant groups are small, effective markets may be impossible to create, simply because of the absence of sufficient numbers of investors. A certain number is necessary to ensure that prediction markets have enough information to aggregate. Nonetheless, government agencies might well enlist such markets to resolve a number of questions, and ambitious efforts are underway to examine how government might enlist them to answer an array of disputed questions.⁷⁹

In fact governments might use prediction markets to help make projections about insolvency, budget deficits, and the costs and benefits of proposed regulations.⁸⁰ In each of these cases, the predictions of prediction markets might provide a “reality check” for deliberative processes. Officials might take into account the markets’ predictions of the anticipated damage from a natural disaster, the number of annual deaths from an actual or anticipated disease (such as mad cow disease or AIDS), the number of American casualties from a war effort, the existence of demonstrable harms from global warming by, say, 2010, the likelihood of scarcity of natural resources, shrinkage of tropical forests in the world, demonstrable deterrent effects from capital punishment or other severe punishments, increases or decreases in emissions of specified air pollutants, increases or decreases in concentrations of air pollution in the ambient air, and much more. In all these cases, private or public institutions might create markets to provide information on crucial questions, and public institutions might take that information into account in making judgments about policy.

The broadest point is that, even when prediction markets are not feasible, an understanding of their virtues helps illuminate the virtues and vices of deliberation—and helps show how to obtain more of the former and less of the latter. Such markets

⁷⁹ See Robert W. Hahn and Paul C. Tetlock, “Using Information Markets to Improve Decision Making,” *Harvard Journal of Law and Public Policy* 29, No. 1 (Fall 2005): 213-289.

⁸⁰ Abramowicz, “Information Markets,” 933.

overcome the collective action problem from which deliberating groups suffer; they also give people a strong incentive to say what they know and to back their best-grounded convictions with money. It should be possible for deliberating groups to learn from the successes of markets, above all by encouraging their members to disclose their privately held information. When such groups do poorly, it is often because they fail to elicit the information that their members have. Good norms, and good incentives, can go a long way toward reducing this problem. Consider here a fundamental redefinition of what it means to be a “team player.” Frequently a team player is thought to be someone who does not upset the group’s consensus. But it would be possible, and a great deal better, to understand team players as those who increase the likelihood that the team will be right—if necessary, by disrupting the conventional wisdom.

The point applies to many organizations, including corporate boards. In the United States, the highest-performing companies tend to have “extremely contentious boards that regard dissent as an obligation” and that “have a good fight now and then.”⁸¹ Investment clubs have little dissent, and lose a great deal of money, when members are united by close social ties.⁸² By contrast, the best-performing investment clubs lack such ties and benefit from dissent and epistemic diversity. When deliberating groups do badly, fear of social sanctions is often a major reason. When they do well, they resemble prediction markets in the sense that their members have a strong incentive to disclose their private information.

Conclusion

Groups often hold a great deal of information, and an important task is to elicit and use the information of their members. Deliberation is generally thought to be the best way of carrying out that task, but deliberative bodies are subject to serious problems. Much of the time, informational influences and social pressures lead members not to say what they know. As a consequence, groups tend to propagate and even amplify cognitive errors. They also emphasize shared information at the expense of unshared information, resulting in hidden profiles. Cascade effects and group polarization are common.

⁸¹ See Jeffrey A. Sonnenfeld, What Makes Great Boards Great, Harvard Business Review (Sept. 2002).

⁸² See Brooke Harrington, Pop Finance: Investment Clubs and the New Ownership Society (Princeton: Princeton University Press, 2006, forthcoming).

Prediction markets have significant advantages over deliberative processes, and in many contexts they might supplement or even replace those processes. Such markets tend to correct rather than amplify individual errors, above all because they allow shrewd investors to take advantage of the mistakes made by others. By providing economic rewards for correct individual answers, they encourage investors to disclose the information they have. As a result, they are often more accurate than the judgments of deliberating groups. To the extent feasible, many groups would often do well to enlist prediction markets in arriving at their judgments, above all because of the accuracy of the price signal. Much more broadly, deliberating groups might attempt to counteract the pressures I have explored, learning from the successes of prediction markets to reduce the risks of deliberative failure.

Readers with comments should address them to:

Professor Cass Sunstein
University of Chicago Law School
1111 East 60th Street
Chicago, IL 60637
csunstei@uchicago.edu

The University of Chicago Law School
Public Law and Legal Theory Working Paper Series

1. Cass R. Sunstein and Edna Ullmann-Margalit, Second-Order Decisions (November 1999; *Ethics*, v.110, no. 1)
2. Joseph Isenbergh, Impeachment and Presidential Immunity from Judicial Process (November 1999; forthcoming *Yale Law and Policy Review* v.18 #1).
3. Cass R. Sunstein, Is the Clean Air Act Unconstitutional? (August 1999; *Michigan Law Review* #3).
4. Elizabeth Garrett, The Law and Economics of “Informed Voter” Ballot Notations (November 1999, *University of Virginia Law Review*, v. 85).
5. David A. Strauss, Do Constitutional Amendments Matter? (November 1999)
6. Cass R. Sunstein, Standing for Animals (November 1999)
7. Cass R. Sunstein, Culture and Government Money: A Guide for the Perplexed (April 2000).
8. Emily Buss, Without Peers? The Blind Spot in the Debate over How to Allocate Educational Control between Parent and State (April 2000).
9. David A. Strauss, Common Law, Common Ground, and Jefferson’s Principle (June 2000).
10. Curtis A. Bradley and Jack L. Goldsmith, Treaties, Human Rights, and Conditional Consent (May 2000; *Pennsylvania Law Review* v. 149).
11. Mary Ann Case, Lessons for the Future of Affirmative Action from the Past of the Religion Clauses? (May 2001, *Supreme Court Review*, 2000)
12. Cass R. Sunstein, Social and Economic Rights? Lessons from South Africa (May, 2000).
13. Jill Elaine Hasday, Parenthood Divided: A Legal History of the Bifurcated Law of Parental Relations (June 2001)
14. Elizabeth Garrett, Institutional Lessons from the 2000 Presidential Election (May 2001).
15. Richard A. Epstein, The Allocation of the Commons: Parking and Stopping on the Commons (August 2001).
16. Jack Goldsmith, The Internet and the Legitimacy of Remote Cross-Border Searches (October 2001).
17. Adrian Vermeule, Does Commerce Clause Review Have Perverse Effects? (October 2001).
18. Cass R. Sunstein, Of Artificial Intelligence and Legal Reasoning (November 2001).
19. Elizabeth Garrett, The Future of Campaign Finance Reform Laws in the Courts and in Congress, The William J. Brennan Lecture in Constitutional Law (December 2001).
20. Julie Roin, Taxation without Coordination (March 2002).
21. Geoffrey R. Stone, Above the Law: Research Methods, Ethics, and the Law of Privilege (March 2002; forthcoming *J. Sociological Methodology* 2002).
22. Cass R. Sunstein, Is There a Constitutional Right to Clone? (March 2002).
23. Emily Buss, Parental Rights (May 2002, forthcoming *Virginia Law Review*).
24. David A. Strauss, Must Like Cases Be Treated Alike? (May 2002).
25. David A. Strauss, The Common Law Genius of the Warren Court (May 2002).
26. Jack Goldsmith and Ryan Goodman, U.S. Civil Litigation and International Terrorism (June 2002).
27. Jack Goldsmith and Cass R. Sunstein, Military Tribunals and Legal Culture: What a Difference Sixty Years Makes (June 2002).
28. Cass R. Sunstein and Adrian Vermeule, Interpretation and Institutions (July 2002).
29. Elizabeth Garrett, Is the Party Over? The Court and the Political Process (August 2002).
30. Cass R. Sunstein, The Rights of Animals: A Very Short Primer (August 2002).
31. Joseph Isenbergh, Activists Vote Twice (November 2002).
32. Julie Roin, Truth in Government: Beyond the Tax Expenditure Budget (November 2002).
33. Cass R. Sunstein, Hazardous Heuristics (November 2002).

34. Cass R. Sunstein, Conformity and Dissent (November 2002).
35. Jill Elaine Hasday, The Principle and Practice of Women's "Full Citizenship": A Case Study of Sex-Segregated Public Education (December 2002).
36. Cass R. Sunstein, Why Does the American Constitution Lack Social and Economic Guarantees? (January 2003).
37. Adrian Vermeule, *Mead* in the Trenches (January 2003).
38. Cass R. Sunstein, Beyond the Precautionary Principle (January 2003).
39. Adrian Vermeule, The Constitutional Law of Congressional Procedure (February 2003).
40. Eric A. Posner and Adrian Vermeule, Transitional Justice as Ordinary Justice (March 2003).
41. Emily Buss, Children's Associational Rights? Why Less Is More (March 2003)
42. Emily Buss, The Speech Enhancing Effect of Internet Regulation (March 2003)
43. Cass R. Sunstein and Richard H. Thaler, Libertarian Paternalism Is Not an Oxymoron (May 2003)
44. Elizabeth Garrett, Legislating *Chevron* (April 2003)
45. Eric A. Posner, Transfer Regulations and Cost-Effectiveness Analysis (April 2003)
46. Mary Ann Case, Developing a Taste for Not Being Discriminated Against (May 2003)
47. Saul Levmore and Kyle Logue, Insuring against Terrorism—and Crime (June 2003)
48. Eric Posner and Adrian Vermeule, Accommodating Emergencies (September 2003)
49. Adrian Vermeule, The Judiciary Is a They, Not an It: Two Fallacies of Interpretive Theory (September 2003)
50. Cass R. Sunstein, Ideological Voting on Federal Courts of Appeals: A Preliminary Investigation (September 2003)
51. Bernard E. Harcourt, Rethinking Racial Profiling: A Critique of the Economics, Civil Liberties, and Constitutional Literature, and of Criminal Profiling More Generally (November 2003)
52. Jenia Iontcheva, Nationalizing International Criminal Law: The International Criminal Court As a Roving Mixed Court (January 2004)
53. Lior Jacob Strahilevitz, The Right to Destroy (January 2004)
54. Adrian Vermeule, Submajority Rules (in Legislatures and Elsewhere) (January 2004)
55. Jide Nzelibe, The Credibility Imperative: The Political Dynamics of Retaliation in the World Trade Organization's Dispute Resolution Mechanism (January 2004)
56. Catharine A. MacKinnon, Directions in Sexual Harassment Law: Afterword (January 2004)
57. Cass R. Sunstein, Black on Brown (February 2004)
58. Elizabeth F. Emens, Monogamy's Law: Compulsory Monogamy and Polyamorous Existence (February 2004)
59. Bernard E. Harcourt, You Are Entering a Gay- and Lesbian-Free Zone: On the Radical Dissents of Justice Scalia and Other (Post-) Queers (February 2004)
60. Adrian Vermeule, Selection Effects in Constitutional Law (March 2004)
61. Derek Jinks and David Sloss, Is the President Bound by the Geneva Conventions? (July 2004)
62. Derek Jinks and Ryan Goodman, How to Influence States: Socialization and International Human Rights Law (March 2004)
63. Eric A. Posner and Alan O. Sykes, Optimal War and *Jus Ad Bellum* (April 2004)
64. Derek Jinks, Protective Parity and the Law of War (April 2004)
65. Derek Jinks, The Declining Significance of POW Status (April 2004)
66. Bernard E. Harcourt, Unconstitutional Police Searches and Collective Responsibility (June 2004)
67. Bernard E. Harcourt, On Gun Registration, the NRA, Adolf Hitler, and Nazi Gun Laws: Exploding the Gun Culture Wars {A Call to Historians} (June 2004)

68. Jide Nzelibe, The Uniqueness of Foreign Affairs (July 2004)
69. Derek Jinks, Disaggregating “War” (July 2004)
70. Jill Elaine Hasday, Mitigation and the Americans with Disabilities Act (August 2004)
71. Eric A. Posner and Cass R. Sunstein, Dollars and Death (August 2004)
72. Cass R. Sunstein, Group Judgments: Deliberation, Statistical Means, and Information Markets (August 2004)
73. Adrian Vermeule, Constitutional Amendments and the Constitutional Common Law (September 2004)
74. Elizabeth Emens, The Sympathetic Discriminator: Mental Illness and the ADA (September 2004)
75. Adrian Vermeule, Three Strategies of Interpretation (October 2004)
76. Cass R. Sunstein, The Right to Marry (October 2004)
77. Jill Elaine Hasday, The Canon of Family Law (October 2004)
78. Adam M. Samaha, Litigant Sensitivity in First Amendment Law (November 2004)
79. Lior Jacob Strahilevitz, A Social Networks Theory of Privacy (December 2004)
80. Cass R. Sunstein, Minimalism at War (December 2004)
81. Eric A. Posner, The Decline of the International Court of Justice (December 2004)
82. Tim Wu, The Breach Theory of Treaty Enforcement (February 2005, revised March 2005)
83. Adrian Vermeule, Libertarian Panics (February 2005)
84. Eric A. Posner and Adrian Vermeule, Should Coercive Interrogation Be Legal? (March 2005)
85. Cass R. Sunstein and Adrian Vermeule, Is Capital Punishment Morally Required? The Relevance of Life-Life Tradeoffs (March 2005)
86. Adam B. Cox, Partisan Gerrymandering and Disaggregated Redistricting (April 2005)
87. Eric A. Posner, Political Trials in Domestic and International Law (April 2005)
88. Cass R. Sunstein, Irreversible and Catastrophic (April 2005)
89. Adam B. Cox, Partisan Fairness and Redistricting Politics (April 2005, *NYU L. Rev.* 70, #3)
90. Cass R. Sunstein, Administrative Law Goes to War (May 2005, *Harvard L. Rev.*, *forthcoming*)
91. Cass R. Sunstein, Chevron Step Zero (May 2005)
92. Bernard E. Harcourt, Policing L.A.’s Skid Row: Crime and Real Estate Development in Downtown Los Angeles [An Experiment in Real Time] (May 2005)
93. Bernard E. Harcourt and Jens Ludwig, Broken Windows: New Evidence from New York City and a Five-City Social Experiment (May 2005)
94. Bernard E. Harcourt, *Against* Prediction: Sentencing, Policing, and Punishing in an Actuarial Age (May 2005)
95. Philip Hamburger, The New Censorship: Institutional Review Boards (May 2005)
96. Eugene Kontorovich, Disrespecting the “Opinions of Mankind” (June 2005)
97. Tim Wu, Intellectual Property, Innovation, and Decision Architectures (June 2005)
98. Lior Jacob Strahilevitz, Exclusionary Amenities in Residential Commons (July 2005)
99. Cass R. Sunstein, Ranking Law Schools: A Market Test? (July 2005)
100. Mary Anne Case, Pets or Meat (August 2005)
101. Adam Samaha, Executive Exposure: Government Secrets, Constitutional Law, and Platforms for Judicial Intervention (August 2005, revised November 2005)
102. Jason J. Czarnezki and William K. Ford, The Phantom Philosophy? An Empirical Investigation of Legal Interpretation (August 2005)
103. Adrian Vermeule, Absolute Voting Rules (August 2005)
104. Eric A. Posner and Adrian Vermeule, Emergencies and Democratic Failure (August 2005)
105. Adrian Vermeule, Reparations as Rough Justice (September 2005)

106. Arthur J. Jacobson and John P. McCormick, The Business of Business Is Democracy (September 2005)
107. Tracey Meares and Kelsi Brown Corkran, When 2 or 3 Come Together (October 2005)
108. Adrian Vermeule, Political Constraints on Supreme Court Reform (October 2005)
109. Lior Jacob Strahilevitz, Information Asymmetries and the Rights to Exclude (November 2005)
110. Cass R. Sunstein, Fast, Frugal and (Sometimes) Wrong (November 2005)
111. Cass R. Sunstein, Justice Breyer's Democratic Pragmatism (November 2005)
112. Adam M. Samaha, Endorsement Retires: From Religious Symbols to Anti-Sorting Principles (November 2005)
113. Adam M. Samaha, Undue Process: Congressional Referral and Judicial Resistance in the Schiavo Controversy (November 2005)
114. Bernard E. Harcourt, Should We Aggregate Mental Hospitalization and Prison Population Rates in Empirical Research on the Relationship between Incarceration and Crime, Unemployment, Poverty, and Other Social Indicators? On the Continuity of Spatial Exclusion and Confinement in Twentieth Century United States (January 2006)
115. Elizabeth Garrett and Adrian Vermeule, Transparency in the Budget Process (January 2006)
116. Cass R. Sunstein, Burkean Minimalism (January 2006)
117. Stephanos Bibas, Transparency and Participation in Criminal Procedure (February 2006)
118. Douglas G. Lichtman, Captive Audiences and the First Amendment (February 2006)
119. Eric A. Posner and Cass R. Sunstein, The Law of Other States (March 2006)
120. Jeff Leslie and Cass R. Sunstein, Animal Rights without Controversy (March 2006)
121. Adrian Vermeule, The Delegation Lottery (March 2006)
122. Adrian Vermeule, Self-Defeating Proposals: Ackerman on Emergency Powers (March 2006)
123. Bernard E. Harcourt, Muslim Profiles Post 9/11: Is Racial Profiling an Effective Counterterrorist Measure and Does It Violate the Right to Be Free from Discrimination? (March 2006)
124. Christine Jolls and Cass R. Sunstein, The Law of Implicit Bias (April 2006)
125. Lior Strahilevitz, "How's My Driving?" for Everyone (and Everything?) (April 2006)
126. Jack Goldsmith and Eric A. Posner, The New International Law Scholarship (May 2006)
127. Eric A. Posner and John Yoo, International Law and the Rise of China (May 2006)
128. Eric A. Posner and Cass R. Sunstein, Chevronizing Foreign Relations Law (May 2006)
129. Jacob E. Gersen and Adrian Vermeule, Chevron as a Voting Rule (June 2006)
130. Jacob E. Gersen, Temporary Legislation (June 2006)
131. Adam B. Cox, Designing Redistricting Institutions (June 2006)
132. Adam B. Cox, The Temporal Dimension of Voting Rights (July 2006)
133. Curtis A. Bradley and Eric A. Posner, Presidential Signing Statements and Executive Power (July 2006)
134. Cass R. Sunstein, Clear Statement Principles and National Security: *Hamdan* and Beyond (July 2006)
135. Cass R. Sunstein, Of Snakes and Butterflies: A Reply (August 2006)
136. Cass R. Sunstein, Montreal versus Kyoto: A Tale of Two Protocols (September 2006)
137. Douglas Lichtman, Irreparable Benefits (September 2006)
138. Adam M. Samaha, Undue Process (September 2006)
139. Eric A. Posner and Adrian Vermeule, The Credible Executive (September 2006)
140. Adam B. Cox and Eric A. Posner, The Second-Order Structure of Immigration Law (November 2006)
141. Eric A. Posner, The International Protection of Cultural Property: Some Skeptical Observations (November 2006)

142. Bernard E. Harcourt and Jens Ludwig, Reefer Madness: Broken Windows Policing and Misdemeanor Marijuana Arrests in New York City, 1989–2000 (December 2006)
143. Bernard E. Harcourt, Embracing Chance: Post-Modern Meditations on Punishment (December 2006)
144. Cass R. Sunstein, Second-Order Perfectionism (December 2006)
145. Wayne Hsiung and Cass R. Sunstein, Climate Change and Animals (January 2007)
146. Cass R. Sunstein, Deliberating Groups versus Prediction Markets (or Hayek's Challenge to Habermas) (January 2007)