



J O I N T C E N T E R
AEI-BROOKINGS JOINT CENTER FOR REGULATORY STUDIES

**Group Judgments:
Deliberations, Statistical Means, and Information Markets**

Cass R. Sunstein*

**Working Paper 04-17
September 2004**

* Karl N. Llewellyn Distinguished Service of Jurisprudence, Law School and Department of Political Science, University of Chicago. For valuable comments, I am grateful to Bernard Harcourt, Douglas Lichtman, Richard Posner, Adam Samaha, and participants in a work-in-progress lunch at the University of Chicago Law School. Many thanks also to Robert Park for extraordinary research assistance. Thanks above all to Reid Hastie, for a great deal of patient help and tutoring.



J O I N T C E N T E R

AEI-BROOKINGS JOINT CENTER FOR REGULATORY STUDIES

In order to promote public understanding of the impact of regulations on consumers, business, and government, the American Enterprise Institute and the Brookings Institution established the AEI-Brookings Joint Center for Regulatory Studies. The Joint Center's primary purpose is to hold lawmakers and regulators more accountable by providing thoughtful, objective analysis of relevant laws and regulations. Over the past three decades, AEI and Brookings have generated an impressive body of research on regulation. The Joint Center builds on this solid foundation, evaluating the economic impact of laws and regulations and offering constructive suggestions for reforms to enhance productivity and welfare. The views expressed in Joint Center publications are those of the authors and do not necessarily reflect the views of the Joint Center.

ROBERT W. HAHN
Executive Director

ROBERT E. LITAN
Director

COUNCIL OF ACADEMIC ADVISERS

KENNETH J. ARROW
Stanford University

MAUREEN L. CROPPER
University of Maryland

PHILIP K. HOWARD
Covington & Burling

PAUL L. JOSKOW
Massachusetts Institute
of Technology

DONALD KENNEDY
Stanford University

ROGER G. NOLL
Stanford University

GILBERT S. OMENN
University of Michigan

PETER PASSELL
Milken Institute

RICHARD SCHMALENSEE
Massachusetts Institute
of Technology

ROBERT N. STAVINS
Harvard University

CASS R. SUNSTEIN
University of Chicago

W. KIP VISCUSI
Harvard University

All AEI-Brookings Joint Center publications can be found at www.aei-brookings.org

© 2004 by the author. All rights reserved.

Executive Summary

How can groups elicit and aggregate the information held by their individual members? There are three possibilities. Groups might use the statistical mean of individual judgments; they might encourage deliberation; or they might use information markets. In both private and public institutions, deliberation is the standard way of proceeding; but for two reasons, deliberating groups often fail to make good decisions. First, the statements and acts of some group members convey relevant information, and that information often leads other people not to disclose what they know. Second, social pressures, imposed by some group members, often lead other group members to silence themselves because of fear of disapproval and associated harms. Hence deliberation often produces a series of unfortunate results: the propagation of errors, hidden profiles, cascade effects, and group polarization. A variety of steps should be taken to ensure that deliberating groups obtain the information held by their members. Information markets have substantial advantages over group deliberation; such markets count among the most intriguing institutional innovations of the last quarter-century and should be used far more than they now are. These points bear on discussion of normative issues, in which deliberation might also fail to improve group thinking, and in which identifiable reforms could produce better outcomes.

“Increased accuracy is a common justification for using groups, rather than individuals, to make judgments. However, the empirical literature shows that groups excel as judges only under limited conditions. . . . [G]roups performing tasks that involve solutions that are not easily demonstrable tend to perform at the level of their average members.”*

“The presumption that Iraq had active WMD programs was so strong that formalized [Intelligence Community] mechanisms established to challenge assumptions and ‘group think,’ such as ‘red teams,’ ‘devil’s advocacy,’ and other types of alternative or competitive analysis, were not utilized.”†

“Sometimes important forecasts are made in traditional group meetings. This . . . should be avoided because it does not use information efficiently. A structured approach for combining independent forecasts is invariably more accurate.”‡

* See Daniel Gigone and Reid Hastie, Proper Analysis of the Accuracy of Group Judgments, 121 *Psych. Bulletin* 149 (1997).

† Select Committee on Intelligence, United States Senate, Report of the U.S. Intelligence Community’s Prewar Intelligence Assessments on Iraq, Conclusions, at 7.

‡ J. Scott Armstrong, Combining Forecasts, in *Principle of Forecasting* 417, 433 (J. Scott Armstrong ed. 2001).

Group Judgments: Deliberations, Statistical Means, and Information Markets

Cass R. Sunstein

1. Introduction

In the last decades, a great deal of attention has been devoted to deliberative accounts of democracy. Its theoretical foundations have been elaborated in some detail,¹ and increasing attention is being devoted to methods for making democratic processes more deliberative. James Fishkin, for example, has pioneered the idea of the “deliberative opinion poll,” by which people are asked to deliberate together on public issues and to state their judgments only after the deliberative process.² Fishkin and Bruce Ackerman have gone so far as to suggest a new national holiday, Deliberation Day, in which people are asked to congregate in groups in order to discuss and debate important issues of public policy.³ Perhaps the proposal is unrealistic; perhaps citizens as a whole should not be expected to deliberate much in a liberal society.⁴ But even if this is true, leaders in the public and private sphere might be urged to deliberate more than they now do, and many accounts of deliberative democracy emphasize the importance of deliberation by representatives.⁵ In any case countless organizations, both public and private, use deliberation to make their judgments; juries and multimember courts, including federal courts of appeals, are only two examples.

Why, exactly, is deliberation important or even desirable? A central answer must be that deliberation will result in wiser judgments and better outcomes.⁶ 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. 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.⁷

¹ See Jurgen Habermas, *Between Facts and Norms* (1998); Amy Gutmann and Dennis Thompson, *Democracy and Disagreement* (1999).

² See James Fishkin, *The Voice of the People* (2000).

³ See Bruce Ackerman and James Fishkin, *Deliberation Day* (2004).

⁴ See Richard A. Posner, *Law, Pragmatism, and Democracy* (2003).

⁵ See William Bessette, *The Mild Voice of Reason* (1996).

⁶ There are other possibilities, of course. Perhaps deliberation has educative effects or contributes to individual self-development; perhaps it legitimates decisions or increases the likelihood that people will acquiesce in them. See Thomas Tyler, *Why People Obey the Law* (1999).

⁷ See Irving Janis, *Groupthink* (2d ed. 1980).

Without structural protections, both private and public groups are likely to err, not in spite of deliberation but because of it.

Consider, for example, the 2004 report of the Senate Select Committee on Intelligence, which explicitly accused the Central Intelligence Agency (CIA) of groupthink, in which the agency's predisposition to find a serious threat from Iraq led it to fail to explore alternative possibilities or to obtain and use the information that it actually held.⁸ In the Committee's view, the CIA "demonstrated several aspects of group think: examining few alternatives, selective gathering of information, pressure to conform within the group or withhold criticism, and collective rationalization."⁹ Thus the agency showed a "tendency to reject information that contradicted the presumption" that Iraq had weapons of mass destruction.¹⁰ Because of that presumption, the agency failed to use its own formalized methods "to challenge assumptions and 'group think,' such as 'red teams,' 'devil's advocacy,' and other types of alternative or competitive analysis."¹¹ Above all, the Committee's conclusions emphasize the CIA's failure to elicit and aggregate information.

This claim is a remarkable and even uncanny echo of one that followed the 2003 investigation of failures at NASA, stressing that agency's similar failure to elicit competing views, including those based on information held by agency employees.¹² The Columbia Accident Investigation Board explicitly attributed the accident to NASA's unfortunate culture, one that does too little to elicit information. In the Board's words, NASA lacks "checks and balances."¹³ It pressures people to follow a "party line."¹⁴ At NASA, "it is difficult for minority and dissenting opinions to percolate up through the agency's hierarchy"¹⁵ -- even though, the Board contended, effective safety programs require the encouragement of minority opinions and bad news.

One of my major goals here is to compare deliberative processes with two nondeliberative means of aggregating privately held information: statistical means and information markets. As we shall see, statistical means sometimes do at least as well as

⁸ Available at <http://intelligence.senate.gov/>.

⁹ *Id.*, conclusions at 4.

¹⁰ *Id.* at 6.

¹¹ *Id.* at 8.

¹² Report of The Columbia Accident Investigation Board, available at http://www.nasa.gov/columbia/home/CAIB_Vol1.html

¹³ *Id.* at 12.

¹⁴ *Id.* at 102.

deliberation. Information markets, which count among the most intriguing institutional innovations of the last quarter-century, often do far better. To explain why deliberation fails, I explore the consequences of two sets of influences on members of deliberating groups.¹⁶ The first consists of informational influences, by which group members fail to disclose what they know because of deference to the information 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 propagate individual errors; emphasize shared information at the expense of unshared information; fall victim to cascade effects; and tend to end up in a more extreme position in line with the predeliberation tendencies of their members.¹⁷ Even federal judges appear to be 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.¹⁸ Statistical groups and information markets are far less susceptible to the pressures that make deliberating groups err.

Because of those pressures, deliberative processes often fail to achieve their minimal goal of aggregating the information that the relevant deliberators actually have. As we shall see, information markets often outperform both statistical and deliberating groups, simply because they are so effective at pooling information. To keep the analysis simple, my principal focus is not on contested judgments of value but on questions with objectively correct answers. An understanding of how deliberation finds, and fails to find, those answers should have implications for its potential and limitations with respect to normative questions as well. The solution to many such questions depends at least in part on resolution of factual issues; it is difficult to take a stand on proposals to raise the minimum wage, to engage in preemptive war, or to overrule Roe v. Wade¹⁹ without resolving several issues of fact. And even when factual issues

¹⁵ Id. at 183.

¹⁶ I explore these mechanisms from a different direction in Cass R. Sunstein, *Why Societies Need Dissent* (2003), but without attention to statistical groups and information markets, and without focusing on propagation of errors, hidden profiles, and the common knowledge effect, which are major emphases here.

¹⁷ This last possibility is emphasized in Roger Brown, *Social Psychology: The Second Edition* (1985); Sunstein, *supra* note.

¹⁸ See Cass R. Sunstein, David Schkade, and Lisa Ellman, *Ideological Voting on Federal Courts of Appeals: A Preliminary Investigation*, 90 Va. L. Rev. 301 (2004); Cass R. Sunstein and David Schkade, *All the President's Judges* (unpublished manuscript 2004).

¹⁹ 410 US 113 (1973).

are not central, deliberation can, in principle, ensure more sensible judgments.²⁰ Unfortunately, however, the problems posed by informational pressure and social influences apply in normative domains as well as elsewhere. I will therefore offer some suggestions for how groups can overcome those problems through structural reforms.

This Article comes in four parts. Part II explores a nondeliberative method for aggregating privately-held information, one that simply takes the average of predeliberation judgments. The resulting judgments of these “statistical groups” are sometimes remarkably accurate, and they provide a useful benchmark for assessing deliberative judgments.²¹ An understanding of the judgments of statistical groups also provides several clues about the usefulness and limitations of relying on the judgments of groups in general. Part III explores the effects of informational pressures and social influences, with an emphasis on propagation of errors, hidden profiles, cascade effects, and group polarization. Part III also attempts to square some apparently conflicting evidence about the performance of deliberating groups; it investigates the possibility that some groups will do as well as or even better than their best members.

Part IV investigates structural reforms that are intended to ensure that group members reveal what they know – for example, by requiring anonymous statements of beliefs before deliberation begins, by assigning specified roles to participants in deliberation, and by structuring incentives to produce disclosure of privately held information. Part V identifies and compares a novel and exceedingly promising method for aggregating individual judgments: “information markets,” in which people bet on the outcomes of events. Information markets have performed remarkably well in many domains, and they have advantages over both statistical judgments and deliberative judgments. They might well be used as a supplement to or even a replacement for collective deliberation. Part VI briefly discusses how the analysis might apply to normative questions.

2. Statistical Groups

Suppose that there is a question about some disputed issue of fact. How many home runs

²⁰ This is the thesis of Fishkin, *The Voice of the People*, *supra* note.

²¹ These are often described as the judgments of “statisticized groups.” See Irving Lorge et al., *A Survey of Studies Contrasting the Quality of Group Performance and Individual Performance, 1920-1957*, 55 *Psych Bull* 337, 344 (1958).

did Hank Aaron hit? When was Calvin Coolidge elected president? Will a district court decision be reversed on appeal? Does a foreign country pose a serious threat to national security? Is the United States likely to have difficulty in winning a particular war? A great deal of evidence suggests that under certain conditions, a promising way to answer such questions is this: Ask a large number of people and take the mean answer. When the relevant conditions are met, the mean answer, which we might describe as the group's "statistical answer,"²² is often accurate, where accuracy is measured by reference to objectively demonstrable fact.

It is well-known that statistical answers from groups of sufficiently large sizes tend to match the views of population-wide samples.²³ This finding bears on issues as diverse as the use of juries as a measure of community sentiment²⁴ and the remarkable success of Google, the search engine; Google is good at finding what a particular searcher wants because it knows what most searchers want.²⁵ But here the question is what is true, not what populations think. Let us therefore explore how statistical groups perform, partly because the answer is important and illuminating itself, and partly because it provides a useful foundation for the assessment of both deliberating groups and information markets.

Evidence

Many of the studies of statistical groups involve quantitative estimates. Consider a few examples:

1. In an early study, Hazel Knight asked college students to estimate the temperature of a classroom.²⁶ Individual judgments ranged from 60 degrees to eighty-five degrees; the statistical judgment of the group was 72.4 degrees, very close to the actual temperature of 72 degrees. That judgment was better than that of 80% of individual judgments.
2. Judging the number of beans in the jar, the group average is almost always better than that of the vast majority of individual members. In one such experiment, a group of

²² See Janis, *supra* note.

²³ See H.J. Eysenck, *The Validity of Judgments As A Function of Number of Judges*, 25 *J. Exp Psych.* 650 (1939).

²⁴ For evidence and comments, see Cass R. Sunstein, Daniel Kahneman, and David Schkade, *Assessing Punitive Damages*, 107 *Yale L.J.* 2071 (1998).

²⁵ See Sergey Brin and Lawrence Page, *The Anatomy of a Large-Scale Hypertextual Web Search Engine*, available at <http://www-db.stanford.edu/>

- fifty-six students was asked about a jar containing 850 beans; the group estimate was 871, a better guess than all but one of students.²⁷
3. Asking two hundred students to rank items by weight, one experimenter found that the group's estimate was 94 percent accurate -- a figure excelled by only five individuals.²⁸
 4. Asked to rank ten piles of buckshot, each only slightly different in size from the others, the group's guess was 94.5% accurate, far more so than that of almost all group members.²⁹
 5. The British scientist Francis Galton sought to draw lessons about collective intelligence by examining a competition in which contestants attempted to judge the weight of a fat ox at a regional fair in England. The ox weighed 1,198 pounds; the average guess, from the 797 contestants, was 1,197 pounds.³⁰

If these findings can be generalized, many questions might plausibly be answered not deliberately, but simply by asking a large group of people and selecting the average response. Imagine that a large company is attempting to project its sales for the following year. Might it do best to poll its salespeople and to choose the average number on the assumption that it is likely to be correct³¹? Or suppose that a company is deciding whether to hire a new employee. Should it ask relevant personnel, not to deliberate, but about their individual views on whether the employee's performance is likely to meet a certain level? Or suppose that the question is whether a case should be settled. Ought a law firm to poll its lawyers about the expected outcome at trial? Or suppose the question is whether a war effort will go well by some identifiable standard. Should the President poll his advisers and take the median answer? To answer these questions, we have to know why, in the relevant studies, the median judgment is so accurate.

The Condorcet Jury Theorem

²⁶ Id. at 342.

²⁷ See James Surowiecki, *The Wisdom of Crowds* 5 (2004).

²⁸ Kate Gordon, *Group Judgments in the Field of Lifted Weights*, 7 *J Exp Psych* 389 (1924); Kate Gordon, *Further Observations on Group Judgments of Lifted Weights*, 1 *J Psych* 105 (1935-1936).

²⁹ R.S Bruce, *Group Judgments in the Field of Lifted Weights and Visual Discrimination*, 1 *J Psych* 117 (1935-1936).

³⁰ See James Surowiecki, *The Wisdom of Crowds* (2004).

³¹ Some affirmative evidence can be found in J. Scott Armstrong, *Combining Forecasts*, in *Principle of Forecasting*

The accuracy of judgments of statistical groups is best explained by reference to the Condorcet Jury Theorem.³² To see how the Theorem works, suppose that people are answering a common question with two possible answers, one false and one true, and that the average probability that each voter will answer correctly exceeds 50 percent. The Jury Theorem holds that the probability of a correct answer, by a majority of the group, increases toward certainty as the size of the group increases.³³ The importance of the Jury Theorem lies in the demonstration that groups are likely to do better than individuals, and large groups better than small ones, if majority rule is used and if each person is more likely than not to be correct. The last proviso is extremely important. Suppose that each individual in a group is more likely to be wrong than right. If so, the likelihood that the group will decide correctly falls to zero as the size of the group increases.

In the context of statistical judgments, several of Condorcet's stringent and somewhat unrealistic assumptions are met. Indeed, the likelihood that they will be met is higher with statistical groups than with deliberating ones. Condorcet assumed that people would be unaffected by whether their votes would be decisive³⁴; that people would not be affected by one another's votes; and that the probability that one group member would be right would be statistically unrelated to the probability that another group member would be right.³⁵ The first two assumptions plainly hold for statistical groups. People do not know what others are saying and hence they cannot be influenced by a belief that their judgments will make the difference to that of the group. The third assumption may or may not be violated. Those who have similar training, or who work closely together, will be likely to see things in the same way,³⁶ and those involved in statistical groups might well meet these conditions. On the other hand, the Condorcet Jury Theorem has been shown to be robust to violations of this third assumption.³⁷

To see why statistical groups perform well, consider the problems just described and note

417, 433 (J. Scott Armstrong ed. 2001).

³² See William P. Bottom et al., Propagation of Individual Bias Through Group Judgment: Error in the Treatment of Asymmetrically Informative Signals, 25 J Risk and Uncertainty 147 (2002).

³³ The theorem is based on some simple arithmetic. Suppose, for example, that there is a three person group, in which each member has a 67% probability of being right. The probability that a majority vote will produce the correct answer is 74%.

³⁴ Bloom et al., *supra* note, at 153.

³⁵ *Id.*

³⁶ *Id.*

³⁷ William P. Bottom et al., Propagation of Individual Bias Through Group Judgment: Error in the Treatment of Asymmetrically Informative Signals, 25 J Risk and Uncertainty 147 (2002).

that even if everyone in the group is not more than 50% likely to be right, the Theorem's predictions may continue to hold. Suppose, for example, that 60% of people are 51% likely to be right and that 40% of people are 50% likely to be right; or that 45% of people are 40% likely to be right and that 55% of people are 65% likely to be right; or even that 51% of people are 51% likely to be right and that 49% of people are merely 50% likely to be right. Even under these conditions, the likelihood of a correct answer will move toward 100% as the size of the group increases. It will not move as quickly as it would if every group member were highly likely to be right, but it will nonetheless move. We could imagine endless variations on these numbers. The point is that even if a significant percentage of the group is not more likely to be right than wrong, or even if many group members are more likely to be wrong than right, an accurate result, from a sufficiently large group, can be expected.

Of course most of the relevant judgments, in studies of statistical groups, do not involve a binary choice; consider the question how many beans are in a jar, how many pounds a given object weighs, or how well a certain product will sell in the following year. But the answers to such questions are not analytically different from those in binary choices. In answering the relevant questions, each person is effectively being asked to answer a long series of binary questions -- ten beans or a thousand beans, twenty beans or five hundred beans, fifty beans or one hundred beans, and so on. If a sufficiently large group is asked to answer such questions, and if most individual answers will be correct, the mean answer will be highly accurate. Of course the combination of probabilities, for a series of binary results, might mean that things will turn out poorly. If someone is 51% likely to answer each of two questions correctly, the probability that she will answer both questions correctly is only slightly higher than 25%. But with large groups, enough people are likely to make good guesses, on the questions involved in certain quantitative judgments, that the average estimate will have a high degree of accuracy.

But compare a situation in which if only 49% of the group is likely to be correct. If so, the likelihood of a mistake will move toward 100% under the same condition. But for the number of beans in a jar, or the weight of an ox, most people are not wholly at sea. The accuracy of the median judgment, for large groups, is simply an application of the Condorcet Jury Theorem. And in certain circumstances, deliberating groups will act in roughly the same way, aggregating their information to produce remarkably accurate results.³⁸

³⁸ Id.

Errors

In this light, we can identify two situations in which the judgment of a statistical group will be incorrect. The first are those in which group members show a systematic bias. The second are those in which their answers are worse than random. The failures of statistical judgments, in these circumstances, have strong implications for deliberation as well.

Bias

A systematic bias in one or another direction will create serious problems for the group's answers. If, for example, an experimenter "anchors" subjects on a misleading number, the median will almost certainly be wrong. Suppose, for example, that a jar contains 800 jelly beans, and the experimenter happens to say, quietly, "many jars of jelly beans, though not necessarily this one, have 500 jelly beans," or even, "I'm asking this question to 250 people."³⁹ In either case, the low number will likely operate as an anchor,⁴⁰ and people's answers will be systematically biased toward understating the actual number, producing an unreliable mean. One study demonstrates more generally that a group's statistical estimate is likely to be erroneous "when the material is unfamiliar, distorted in a way such that all individuals are prone to make similar errors of estimation."⁴¹ The error-producing effects of anchors are simply a special case of this general point. Anchors are undoubtedly at work in deliberating groups as well, although in theory deliberation might reduce their effects.

Random or worse

Suppose that people are asked not about the number of jelly beans in a jar, but about the number of atoms in a jelly bean. On that question, people's answers are hopelessly ill-informed, and there is no reason at all to trust their judgments. Consider a small-scale study at the University of Chicago Law School, one that strongly supports this conclusion. A number of faculty members were asked the weight, in pounds, of the fuel that powers space shuttles. The actual answer is 4 million pounds. The median response was 200,000; the mean was 55,790,555

³⁹ Even self-evidently arbitrary anchors have significant effects on people's judgments. See Gretchen Chapman and Eric Johnson, *Incorporating the Irrelevant: Anchors in Judgments of Belief and Value*, in *Heuristics and Biases: The Psychology of Intuitive Judgment* 120 (Thomas Gilovich et al. eds. 2002).

⁴⁰ See *id.*

(because of one outlier choice) -- both wildly inaccurate. In a binary choice, of course, people's answers will be worse than random only if they are unaware of how little they know; if they know that they are likely to be wrong, they should choose randomly, which gives them a 50% probability of being right. But sometimes people think they know more than they do, and many tasks do not involve binary choices at all. Statistical groups will err if confusion and ignorance are so widespread that individual's answers are worse than random.

Statistical Answers and Experts

Might statistical means be used more than they now are? Do statistical means outperform experts? Everything depends on the competence of the experts. If we could find real experts on the weight of oxen or on how to count jelly beans, and if we understand expertise to be the ability to make accurate assessments, then experts would by stipulation do better than statistical means. Suppose, for example, that a deliberating group of lawyers is trying to decide how many Supreme Court decisions have invalidated a state or federal law, or the number of lines in Antigone, or the weight of the most recent winner of the Kentucky Derby. Would it make any sense to poll the lawyers individually and to assume that the mean response is accurate? The studies outlined above suggest that if the group is large enough, the mean answer will be at least good.⁴² But there are many ways to do far better.

For many factual questions, of course, a little research would be sufficient to identify the correct answers. But for some factual issues, even significant research is inconclusive, and it is best to consult experts. And if experts are available, it would make sense to obtain a statistical answer from them, rather than to select one or a few. If experts are likely to be right, a statistical group of experts should have the same advantage over individual experts as a statistical group of ordinary people has over ordinary individuals. In fact a great deal of evidence supports this claim.⁴³ In a series of thirty comparisons, statistical groups of experts had 12.5% fewer errors on forecasting tasks involving such diverse issues as company earnings, cattle and chicken prices, real and nominal GNP, survival of patients, and housing starts.⁴⁴ For example, statistical groups

⁴¹ Lorge et al., *supra* note, at 346.

⁴² I conducted such a poll with faculty at the University of Chicago Law School, who did fairly well in estimating the weight of the horse who won the Kentucky Derby, fairly badly in estimating the number of lines in Antigone -- and horrendously with the number of Supreme Court invalidations of state and federal law!

⁴³ See J. Scott Armstrong, *Combining Forecasts*, in *Principles of Forecasting* 416 (2001).

⁴⁴ *Id.* at 428.

of experts significantly outperformed individual experts in predicting the annual earnings of firms; changes in the American economy; and annual peak rainfall runoff in eight different countries.⁴⁵ The implication is straightforward: “Organizations often call on the best expert they can find to make important forecasts. They should avoid this practice, and instead combine forecasts from a number of experts.”⁴⁶

Consider in this regard the Copenhagen Consensus, designed to inform policy judgments about global risks.⁴⁷ The Copenhagen Consensus was obtained for a series of possible interventions, involving climate change, water and sanitation, hunger and malnutrition, free trade, and communicable diseases, among others. A number of experts were asked about the best way to promote global welfare, and particularly the welfare of developing countries, assuming that \$50 billion were made available. The experts ranked the possible projects, producing an overall ranking (reflecting the mean rankings of the experts taken as a whole).⁴⁸ I do not mean to suggest that the results of this particular exercise are correct; everything depends on whether the relevant experts were in a position to offer reliable answers on the questions at hand. But if statistical means are a good way to aggregate knowledge when ordinary people know something of relevance, then they are also a good way to aggregate knowledge from experts.

3. Deliberating Groups

Although the judgments of statistical groups can be quite accurate, it is easy to imagine that a deliberating group would be much better. In principle, a deliberating group should do well even when its members are error-prone. Deliberation, in the form of an exchange of information and reasons, might well bring them into line. If many group members give answers that are worse than random, perhaps other group members can show them how they have erred. If individuals have been manipulated in their private judgments, perhaps deliberation will undo the effects of the manipulation.⁴⁹ If individual members have anchored on a misleading value, perhaps deliberation will expose the anchor as such.

To make the analysis tractable, let us focus on how deliberating groups might be able to

⁴⁵ Id. at 430-31.

⁴⁶ Id. at 433.

⁴⁷ See <http://www.copenhagenconsensus.com/>.

⁴⁸ See id.

solve factual questions or cognitive puzzles that have correct solutions. The latter are often questions of instrumental rationality, posing a question about the right strategy for achieving agreed-upon goals.

Mechanisms and Realities

If groups perform better than their average member, we can imagine three principal mechanisms by which the improvement occurs.

- Groups as equivalent to their best members. One or more group members will often know the right answer, and other group 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 some form of bias that leads to error, other group members might correct them. Suppose, for example, that a panel of judges is trying to recall relevant Supreme Court decisions in a somewhat specialized area. If one of the judges is actually aware of those decisions, the group will be made aware of them too. Or suppose that a group of military officials is attempting to assess the strengths and weaknesses of a potential enemy in some part of the world. If one of them is a specialist, all of them can learn what the specialist knows. Many deliberating groups contain at least one expert on the question at hand; if group members listen to the expert, they will do at least as well as she does.
- The whole as 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 group members, so that the group is potentially expert even if its members are not. Deliberation might elicit the relevant information and allow the group to make a sensible judgment. Almost everyone has had the experience of being a part of a group that ended up with a solution that went beyond what any individual member could

⁴⁹ See Ackerman and Fishkin, *supra* note.

have produced on her own. In this process, the whole is equal to the sum of the parts – and the sum of the parts is what is sought.

- 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 their 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. And in fact, groups sometimes do outperform their best members.⁵⁰

To what extent do these mechanisms work in practice? Two points are entirely clear. First, deliberation usually reduces variance.⁵¹ After talking with together, group members tend to come into accord with one another.⁵² Statistical groups thus show far more diversity of opinion than deliberating groups. Second, group members tend to become far more confident of their judgments after they speak with one another.⁵³ A significant effect of group interactions is a greater sense that one's post-deliberation conclusion is correct – whether or not those interactions actually increase accuracy. After people have deliberated with one another, they are highly likely to have a heightened sense that their view is right. Corroboration by others increases confidence in one's judgments.⁵⁴ It follows that that members of deliberating groups will usually converge on a position on which group 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 situation).

Unfortunately, there is no systematic evidence that deliberating groups will usually perform well. With respect to questions with definite answers, deliberating groups tend to do about as well as or slightly better than their average member, but not as well as their best members.⁵⁵ Hence it is false to say that group members usually end up deferring to their internal

⁵⁰ See Gigone and Hastie, *supra* note.

⁵¹ See Roger Brown, *Social Psychology: The Second Edition* 206-07 (1985).

⁵² *Id.*

⁵³ See Chip Heath and Rich Gonzalez, *Interaction With Others Increases Decision Confidence But Not Decision Quality: Evidence against Information Collection Views of Interactive Decision Making*, 61 *Org Behavior and Human Decision processes* 305 (1995).

⁵⁴ See Robert Baron et al., *Social Corroboration and Opinion Extremity*, 32 *J Experimental Soc. Psych.* 537 (1996).

⁵⁵ See Gigone and Hastie, *supra* note; Reid Hastie, *Experimental Evidence of Group Accuracy*, in *Information*

specialists. No significant differences are found between deliberating groups and average individual performances in numerical estimates, such as assessment of the number of beans in a jar or the length of lines.⁵⁶ One study finds that when asked to estimate the populations of American cities, groups did as well as their most accurate individual member⁵⁷; but this is an atypical result.⁵⁸ Another study attempted to test whether deliberating groups were particularly good at telling whether people were telling the truth or instead lying.⁵⁹ The individual votes, predeliberation, were 48 percent correct, about the same as the post-deliberation judgments. Approximately the same number of people shifted toward error as toward correct answers.

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 as well.⁶⁰ With experts, the same general conclusion holds. A “structured approach for combining independent forecasts is invariably more accurate” than “traditional group meetings,” which do “not use information efficiently.”⁶¹

Let us discuss the key sources of deliberative failure, understood as a failure to make good decisions on the basis of the information that group members actually have.

Two Sources of Deliberative Failure: Informational Influences and Social Pressures

A primary advantage of statistical groups is that members say what they think. But with deliberating groups, this might not happen. Exposure to the views of others might lead people to silence themselves, and for two different reasons.

The first reason 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 that 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

Pooling and Group Decision Making 129 (Bernard Grofman and Guillermo Owen et al. eds. 1983).

⁵⁶ Id. at 133.

⁵⁷ Hillel Einhorn et al., Quality of Group Judgment, 84 *Psych Bulletin* 158 (1977).

⁵⁸ See Hastie, *supra* note.

⁵⁹ See id.

⁶⁰ Robert J. MacCoun, Comparing Micro and Macro Rationality, in *Judgments, Decisions, and Public Policy* (M.V. Gowda and Jeffrey Fox eds. 2002).

⁶¹ Armstrong, *supra* note, at 433.

statements of others.⁶² 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 the sole dissenter. If all but one person in a deliberating group has said that X is true, then the remaining member is likely to agree that X is true, even to the point of ignoring the evidence of his own senses.⁶³ 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.⁶⁴

Informational signals come in three different forms, involving conduct, conclusions, and reason-giving. First, group members might purchase certain products, visit particular places, or engage in certain actions; their conduct will provide a signal about their beliefs. Second, group members might express their conclusions about some issue. They might say that global warming is a serious problem, that crime is rising in New York City, that minimum wage legislation increases unemployment. Third, group members might give reasons and arguments for their beliefs, going beyond conclusions to explain why they think as they do. If a number of different arguments favor a certain conclusion, and if each of these arguments is plausible, there is more reason to think that the conclusion is right. Conduct, conclusions, and reasons will have different effects in different circumstances; we can imagine a group whose members are unimpressed by conclusions but much affected by behavior, or a group whose members pay far more attention to reasons than to conclusions.⁶⁵ By definition, the deliberative ideal is supposed to include reason-giving, not merely actions or statements of conclusions.

The second reason involves social influences. If people fear that their statements will be disliked or ridiculed, they might not speak out, even on questions of fact. Their silence might stem not from a belief that they are wrong, as in the case of informational pressure, but instead but 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

⁶² Cf. Andrew Caplin & John Leahy, *Miracle on Sixth Avenue: Information Externalities and Search*, 108 *Econ. J.* 60, 61 (1998).

⁶³ See the overview in Solomon Asch, *Opinions and Social Pressure*, in *Readings About the Social Animal* 13 (Elliott Aronson ed.) (New York: W.H. Freeman, 1995).

⁶⁴ See David Krech et al., *Individual in Society* 514 (New York: McGraw Hill, 1962)

⁶⁵ For relevant data, see Gene Rowe and George Wright, *Experts Opinions in Forecasting: The Role of the Delphi*

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 defy them publicly.

Both informational pressure and social influences help explain the finding that in a deliberating group, those with a minority position often silence themselves or otherwise have disproportionately little weight.⁶⁶ There is a more particular finding: Members of low status groups – less educated people, African-Americans, sometimes women – carry less influence within deliberating groups than their higher-status peers.⁶⁷ The unfortunate result can be a loss of information to the group as a whole, in a way that ensures that deliberating groups do far less well than they would if only they could aggregate the information held by group members.

Informational pressure and social pressures also help explain some otherwise puzzling findings about judicial voting on federal courts of appeals. Consider the fact that on three-judge panels, Republican appointees show far more conservative voting patterns when sitting with two other Republican appointees – and that Democratic appointees show far more liberal voting patterns when sitting with two other Democratic appointees.⁶⁸ Consider too the finding that when sitting with two Republican appointees, Democratic appointees show quite conservative voting patterns, close to those of Republican appointees in the aggregate data – and that when sitting with two Democratic appointees, Republican appointees are fairly liberal, with overall votes akin to those of Democratic appointees.⁶⁹ Informational pressure and social influences are not the whole story, but they play a substantial role.⁷⁰

More generally, a comprehensive study demonstrates that majority pressures can be powerful even for factual questions on which some people know the right answer.⁷¹ Both informational pressure and social influences help account for the power of those influences. The study involved 1200 people, forming groups of six, five, and four members. Individuals were

Technique, in *Principles of Forecasting* 125 (J. Scott Armstrong ed. 2001).

⁶⁶ See Glenn Loury, *Self-Censorship in Public Discourse: A Theory of “Political Correctness” and Related Phenomena*, *Rationality and Society* 428 (1994).

⁶⁷ See Caryn Christenson and Ann Abbott, *Team Medical Decision Making*, in *Decision Making in Health Care* (Gretchen Chapman and Frank Sonnenberg eds.) (New York: Cambridge University Press, 2000), at 267, 273-76.

⁶⁸ See Cass R. Sunstein, David Schkade, and Lisa Ellman, *Ideological Voting on Federal Courts of Appeals: A Preliminary Investigation*, 90 *Va. L. Rev.* 301, 314 (2004); Cass R. Sunstein and David Schkade, *All the President’s Judges* (unpublished manuscript 2004).

⁶⁹ *Id.*

⁷⁰ *Id.* at 337-46.

⁷¹ R.L. Thorndike, *The Effect of Discussion Upon the Correctness of Group Decisions, When the Factor of Majority Influence Is Allowed For*, 9 *J. Social Psych.* 343 (1938).

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 the group's answers. The truth played a role too, but a lesser one. If a majority of individuals on the group gave the right answer, the group's majority moved toward the majority in 79 percent of the cases. If a majority of individuals on the group gave the wrong answer, the group 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 vs. 62 percent). What is important is that groups did not take perform as well as they would have if they had properly aggregated the information that group members had.

Informational influences and social pressures lead to four kinds of failures on the part of deliberating groups. Each has a distinctive structure; I discuss them in sequence.

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 are also productive of errors.⁷² A growing literature explores the role of these heuristics and biases⁷³ and their relationship to law and policy.⁷⁴ 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).⁷⁵ The representative heuristic helps explain what Paul Rozin and Carol Nemeroff call “sympathetic magical thinking,” including the beliefs that some objects have contagious properties and that causes resemble their effects.⁷⁶ People also err because they use the availability heuristic to answer difficult questions about probability. When people use this heuristic, they answer a question of

⁷² For an overview, see *Heuristics and Biases: The Psychology of Intuitive Judgment* (Thomas Gilovich et al. eds. 2002).

⁷³ See *id.*

⁷⁴ See *Behavioral Law and Economics* (Cass R. Sunstein ed. 2000).

⁷⁵ Amos Tversky and Daniel Kahneman, Judgment under Uncertainty: Heuristics and Biases 3, in *Heuristics and Biases: The Psychology of Intuitive Judgment* (Thomas Gilovich et al. eds, 2002)

⁷⁶ Paul Rozin and Carol Nemeroff, Sympathetic Magical Thinking: The Contagion and Similarity “Heuristics,” *id.* at 201.

probability by asking whether examples come readily to mind.⁷⁷ People are strikingly vulnerable to framing effects, making different decisions depending on the wording of the problem. For a simple example, consider the question whether to undergo a risky medical procedure. When people are told, “Of those who have this procedure, 90 percent are alive after five years,” they are far more likely to agree to the procedure than when they are told, “Of those who have this procedure, 10 percent are dead after five years.”⁷⁸

For purposes of assessing deliberation, a central question is whether groups avoid the errors of the individuals who compose them. There is no clear evidence that they do, and there is considerable evidence that 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.

Consider some key findings. Groups have been found to amplify, rather than to attenuate, reliance on the representativeness heuristic⁷⁹; to reflect even larger framing effects than individuals⁸⁰; to show more overconfidence than group members⁸¹; to be more affected by the biasing effect of spurious arguments from lawyers⁸²; to be more susceptible to the “sunk cost fallacy”⁸³; and to be more subject to choice-rank preference reversals.⁸⁴ 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.⁸⁵ In addition, groups demonstrate essentially the same level of reliance on the availability heuristic, even when use of that heuristic leads to clear errors.⁸⁶

Why are cognitive errors propagated and often amplified at the group level?

⁷⁷ See Tversky and Kahneman, *supra* note 2, at 3, 11-14.

⁷⁸ See Donald Redelmeier, Paul Rozin, & Daniel Kahneman, Understanding Patients' Decisions, 270 JAMA 72, 73 (1993).

⁷⁹ M.F. Stasson et al., Group Consensus Approaches on Cognitive Bias Tasks, 30 Japanese Psych. Res. 68 (1988).

⁸⁰ See Kerr et al., *supra* note.

⁸¹ Dunning and Ross (1992); J.A. Sniezek and R.A. Henry, Accuracy and Confidence in Group Judgment, 42 Org Behav and Human Decision Processes 1 (1989).

⁸² E.L. Schmann and W.C. Thompson, Effects of Attorney's Arguments on Jurors' Use of Statistical Evidence (unpublished manuscript 1989).

⁸³ G. Whyte, Escalating Commitment in Individual and Group Decision Making, 54 Org Behavior and Human Dec. Processes 430 (1993).

⁸⁴ J.C. Mowen and J. W. Gentry, Investigation of the Preference Reversal Phenomenon in a New Product Introduction Task, 65 J Applied Psych 715 (1980); J.R. Irwin and J.H. Davis. Choice/Matching Preference Reversals in Groups, 64 Org Behavior and Human Dec Processes 325 (1995).

⁸⁵ G. Whyte, Escalating Commitment in Individual and Group Decision Making, 54 Org Behavior and Human Dec. Processes 430 (1993).

⁸⁶ M.F. Stasson et al., Group Consensus Processes on Cognitive Bias Tasks, 30 Japanese Psych. Res. 68 (1988).

Informational pressures and social influences are unquestionably at work. Suppose, for example, that most members of a group are prone to make conjunction errors (believing that A and B are more likely to be true than A alone).⁸⁷ If the majority makes conjunction errors, then most people will see conjunction errors, and what they see will convey information about what is right. Those who are not specialists in logic are likely to think: If most people make conjunction errors, perhaps they are not errors at all. Of course some people will not fall prey to those errors and may even correct them; but group members would have to have a high degree of confidence to do so. Recall here the finding that groups make more conjunction errors than individuals when the rate of individual error is high⁸⁸ – a finding that fits well with the informational explanation of why group amplify errors.

Social influences undoubtedly contribute to the propagation and amplification of individual mistakes. If most group members make conjunction errors, others might make them too simply in order not to seem disagreeable or foolish – at least if there is no particular incentive to produce the right answer. And if most group members use the availability heuristic, or commit the sunk-cost fallacy, then there will be social pressure to do the same.

To be sure, there is some evidence of group attenuation of biases. For example, groups are slightly less susceptible to hindsight bias.⁸⁹ Apparently members who are not susceptible to that bias are able to persuade others that it is indeed a bias. Groups are especially likely to outperform the average individual when members are subject to “egocentric biases.”⁹⁰ When asked what percentage of other undergraduates will vote for George W. Bush, have cell phones, watch television on Tuesday night, enjoy a particular singer, or believe that Spiderman 2 will win at least one Oscar, most people show a bias in the direction that they themselves favor. They believe that their tastes and preferences are typical. But in groups with diverse views, individual members learn that their own position is not universally held, and hence the bias is reduced.⁹¹ Group deliberation supplies an important corrective.

But the more general point is that with group discussion, individual errors are usually

⁸⁷ See Amos Tversky and Daniel Kahneman, Extensional versus Intuitive Reasoning: The Conjunction Fallacy in Probability Judgment, in *Heuristics and Biases: The Psychology of Intuitive Judgment* 19 (Thomas Gilovich et al. eds. 2002).

⁸⁸ *Id.*

⁸⁹ D. Stahlberg et al., We Knew It All Along: Hindsight Bias in Groups, 63 *Org Behavior and Human Decision Processes* 46 (1995).

⁹⁰ Personal communication, Reid Hastie, University of Chicago Business School, who has conducted experiments on this issue for many years.

propagated, not eliminated,⁹² and amplification of mistakes is more likely than alleviation. A general review suggests that when individuals show a high degree of bias, groups are likely to be more biased, not less biased, than their median or average member; in such circumstances, “groups generally can be expected to amplify rather than correct individual bias.”⁹³ This point is an application of the lesson, from the Condorcet Jury Theorem, that as the size of the group expands, the likelihood of group error expands toward 100% if each group member is more likely to be wrong than right. What I am emphasizing here is that social dynamics can aggravate rather than reduce that problem. And if this is so, then jury deliberations, as well as deliberation within multimember courts, are prone to error.

Deliberative Failure, 2: Hidden Profiles and Common Knowledge

Suppose that group members have a great deal of information – enough to produce the 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, many studies demonstrate that this unfortunate result is highly likely. “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.⁹⁴ 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.

Examples

Consider a study of serious errors within working groups, both face-to-face and online.⁹⁵

⁹¹ Id.

⁹² See Bottom et al., *supra* note.

⁹³ Robert J. MacCoun, Comparing Micro and Macro Rationality, in *Judgments, Decisions, and Public Policy* (M.V. Gowda and Jeffrey Fox eds. 2002).

⁹⁴ Daniel Gigone and Reid Hastie, The Common Knowledge Effect: Information Sharing and Group Judgments, 65 *J Personality and Social Psych* 959 (1993).

⁹⁵ See R. Hightower and L. Sayeed, The Impact of Computer-Mediated Communication Systems on Biased Group Discussion, 11 *Computers in Human Behavior* 33 (1995).

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 on-line. 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.”⁹⁶

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.⁹⁷ In one condition, each member of the four-person groups was given most of the relevant information (66% of the information about each candidate). In that condition, 67% of group members favored Candidate A before discussion, and 85% after discussion.⁹⁸ This is 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% of information about each candidate was shared, and 67% was unshared. As the condition was designed, the shared information favored unambiguously inferior candidates, B and C; but if the unshared information emerged through discussion, and was taken seriously, Candidate A would be chosen. In that condition, less than 25% of group members favored Candidate A before discussion, a natural

⁹⁶ Patricia Wallace, *The Psychology of the Internet* 82 (1999).

⁹⁷ See Garold Stasser and William Titus, *Pooling of Unshared Information in Group Decision Making: Biased Information Sampling During Discussion*, 48 *J Per and Social Psych* 1467 (1985).

⁹⁸ *Id.* at 1473. See also Garold Stasser and William Titus, *Hidden Profiles: A Brief History*, 14 *Psych Inquiry* 304

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.⁹⁹ In other words, groups did worse, not better, than individuals when 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.”¹⁰⁰ That conclusion has a clear connection with the judgments, mentioned above, about large-scale information failures at the CIA and similar failures at NASA.¹⁰¹ It follows that “[g]roup 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.”¹⁰² 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.¹⁰³ 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.”¹⁰⁴

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 the most substantial influence on group judgments, far more than information held by one member or a few.¹⁰⁵ 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 that item before the group discussion and judgment.”¹⁰⁶

(2003).

⁹⁹ Id.

¹⁰⁰ Id. at 1476.

¹⁰¹ See supra.

¹⁰² Id.

¹⁰³ Id.

¹⁰⁴ Stasser and Titus, *Hidden Profiles*, at 305.

¹⁰⁵ See Daniel Gigone and Reid Hastie, *The Common Knowledge Effect: Information Sharing and Group Judgments*, 65 *J Personality and Social Psych* 959 (1993).

¹⁰⁶ Id. at 960.

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.”¹⁰⁷

In a key study, deliberating groups would have lost nothing in terms of accuracy if they had simply averaged the judgments of the people involved – a clear finding that deliberation may not improve on the judgments of statistical group.¹⁰⁸ The more shared information is (the more that it stands as “common knowledge”), the more impact it will have on group members before discussion begins – and the more impact it will have as discussion proceeds, precisely because commonly held information is more likely to be discussed.

As might be expected, the group’s focus on shared information increases with the size of the group.¹⁰⁹ In another study designed to test judgments about candidates for office, involving both three-person and six-person groups, all discussions focused far more on shared information than on unshared information – but the effect was significantly greater for six-person groups. Most remarkably, “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.”¹¹⁰ And despite the failures of their deliberations, group members were significantly more confident in their judgments after discussion.¹¹¹

How can these findings be squared with the Condorcet Jury Theorem? The most fundamental point is that in deliberation, individuals are not making judgments on their own; they are being influenced by the judgments of others. When interdependent judgments are being made, and when some people are wrong, the Condorcet Jury Theorem offers no clear predictions. Under such circumstances, it is not at all clear that groups will do better than individuals.¹¹² And when groups fail, the tendency toward hidden profiles is often part of the reason.

Informational influences and social pressures redux

¹⁰⁷ Id. at 973.

¹⁰⁸ Id.

¹⁰⁹ See Garold Stasser et al., Information Sampling in Structured and Unstructured Discussions of Three and Six-Person Groups, 57 *J Personality and Social Psych* 67 (1989).

¹¹⁰ Id. at 78.

¹¹¹ Id. at 72.

¹¹² For evidence, see Norbert Kerr et al., Bias in Judgment: Comparing Individuals and Groups, 103 *Psych. Rev.* 687 (1996). On some of the theoretical issues, see David Austen-Smith and J.S. Banks, Information Aggregation,

Why do hidden profiles remain hidden? The two major explanations track the informational and social accounts traced above. When information is held by all or many, it is more likely, as a statistical matter, to be repeated in group discussion, and hence more likely to be influential than information that is held by one person or a few.¹¹³ There are two different points here.¹¹⁴ Information held by all or most group members is likely to influence individual judgments, and those judgments will in turn affect the judgments of the group. Thus the effects of a shared piece of information will influence the group simply through its impact on predeliberation judgments. In addition, shared information, simply because it is shared, is more likely to be explored during group discussion. Suppose, for example, that a team of five lawyers is deciding whether to appeal an adverse trial court ruling. If each of the five lawyers shares certain information indicating that an appeal would be unsuccessful, that information is more likely to emerge in group discussion than separate parcels of information, individually held by each lawyer, suggesting that an appeal would likely succeed. If the team of lawyers stresses the information that is held by all, that information will have a disproportionate influence on its ultimate decision.¹¹⁵ This is a statistical point about information sampling.

But information sampling provides an incomplete account; hidden profiles are even more hidden than would be predicted by that account.¹¹⁶ To understand the additional element, consider the finding that low-status members of groups are “increasingly reluctant over the course of discussion to repeat unique information.”¹¹⁷ Those in a group who are inexperienced, or are thought to be low on the hierarchy, are particularly loathe to emphasize their privately held information as discussion proceeds. This finding suggests that group members, and especially lower status ones, are alert to the reputational costs of emphasizing information that most group members seem to lack. Lower status members “are likely to drop unique information like a hot potato” -- partly because of the difficulty of establishing its credibility and relevance,¹¹⁸ partly because they may incur group disapproval if they press a line of argument that others reject. It follows that hidden profiles are produced by both informational and reputational pressures imposed by the initial distribution of views.

Rationality, and the Condorcet Jury Theorem, 90 *American Political Science Review* 34 (1996).

¹¹³ See Stasser and Titus, *Hidden Profiles*, *supra* note, at 306-07.

¹¹⁴ See Gigone and Hastie, *supra*, at 960.

¹¹⁵ *Id.*

¹¹⁶ Stasser and Titus, *supra*, at 308.

¹¹⁷ *Id.*

In the same vein, those who discuss shared information obtain rewards in the form of an enhanced sense of competence and credibility in the eyes of others – and in their own eyes as well.¹¹⁹ In both face-to-face discussions and purely written tasks, people give higher ratings (in terms of knowledge, competence, and credibility) to themselves and to others after receiving information that they knew already. It follows that “a bearer of valuable, unshared information may need to establish credibility by telling others what they already know before telling them what they do not already know.”¹²⁰ The general problem is that deliberating groups often perform poorly because they fail to elicit information that could steer them in the right directions.

Deliberative Failure, 3: Cascades

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, suppose that doctors are deciding whether to prescribe a specified therapy for menopausal women. Let us assume that if the specified therapy creates significant risks of heart disease, it should not be prescribed; if it does not create such risks, it is worthwhile.¹²¹ Let us also assume that the doctors are deciding in sequence, in some kind of temporal queue, and each doctor knows his place on that queue. From his own experience, each doctor has some private information about what should be done. But each doctor also cares, reasonably enough, about the judgments of others. Anderson is the first to decide. He prescribes the specified therapy if he thinks that the risk is low but declines if he thinks that the risk is high. Suppose that Anderson prescribes. Barber now knows Anderson’s judgment; it is clear that she too should certainly urge the specified therapy if she makes the same judgment independently. But if her independent judgment is that the risk is high, she would

¹¹⁸ Id.

¹¹⁹ See G.M. Wittenbaum et al., Mutual Enhancement: Toward an Understanding of the Collective Preference for Shared Information, 77 *J Pers and Social Psych* 967 (1999).

¹²⁰ Stasser and Titus, *supra*, at 311.

¹²¹ I draw here on David Hirschleifer, The Blind Leading the Blind, in *The New Economics of Human Behavior* 188, 193-94 (Marianno Tommasi and Kathryn Ierulli eds.) (1995).

-- if she trusts Anderson no more and no less than she trusts herself -- be indifferent about whether to prescribe, and might simply flip a coin.

Now turn to a third doctor, Carlton. Suppose that both Anderson and Barber have prescribed the specified therapy, but that Carlton's own information, though inconclusive, suggests that the risk is probably high. In that event, Carlton might well ignore what he knows and prescribe the therapy. It is likely, in these circumstances, that both Anderson and Barber saw a low risk, 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 Anderson and Barber did, not on the basis of his own information, and also that subsequent doctors know what Anderson, Barber, and Carlton did. On reasonable assumptions, they will do exactly what Carlton did: prescribe the specified therapy regardless of their private information (which, we are supposing, is relevant but inconclusive). This will happen even if Anderson blundered.¹²²

If this is what is happening, there is a serious social problem: Doctors who are in the cascade do not disclose, to their successors and to the public, the information that they privately hold. In the example just given, doctors' actions will not reflect the overall knowledge, or the aggregate knowledge, of the health consequences of the specified therapy -- even if the information held by individual doctors, if actually revealed and aggregated, would give a quite accurate picture of the situation. The reason for the problem is that individual doctors are following the lead of those who came before. Subsequent doctors might fail to rely on, and fail to reveal, private information that actually exceeds the information collectively held by those who started the cascade. Hence the medical profession generally will lack information that both doctors and patients need to have.

The cascade just discussed does not involve group deliberation, but it should be easy to see how cascades can arise during group processes, as individuals give some signal, through behavior, conclusion, or analysis, about what they believe. To be sure, in the standard informational cascade those who follow are responding to actions rather than to explanations or to sustained talk. At first glance, this point seems to suggest that cascades will occur only when people do not talk and exchange reasons. And it is plausible to suggest that bad cascades are less likely when people give reasons. But this point should not be oversold. Within deliberating

¹²² Id. at 195.

groups, some efforts to give reasons are not much more informative than the conclusion itself. And when the explanation is both informative and clear, people might fall into a bad cascade simply because they do not believe that they have enough information to reject it.

Cascades involving group processes 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.¹²³ In each period, the contents of the chosen urn were emptied in a container. A randomly selected subject was asked to make one (and only one) private draw of a ball. After that draw, the subject recorded, on an answer sheet, the color of the draw and her own decision about which urn was involved. The subject did not announce her draw to the group, but she did announce her own decision to everyone. Then the urn was passed to the next subject for her own private draw, which again was not disclosed, and her 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. Subjects could earn \$2 for a correct decision. If the subject picks the urn based only on her private information, she will be right 66.7 percent of the time. The point of the experiment is 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 they often produced errors. After a number of individual judgments were revealed, people sometimes announced decisions that were inconsistent with their private draw, but that fit with the majority of previous announcements.¹²⁴ Over 77% of “rounds” resulted in cascades, and 15% of private announcements did not reveal a “private signal,” that is, the information provided by people’s own draw. 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% of the time – far less than a 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. Notably, the majority of decisions were rationally based

¹²³ See Lisa Anderson and Charles Holt, *Information Cascades in the Laboratory*, 87 *Am Econ Rev* 847 (1997).

¹²⁴ See Angela Hung and Charles Plott, *Information Cascades: Replication and an Extension to Majority Rule and Conformity-Rewarding Institutions*, 91 *Am Econ Rev* 1508, 1515 (2001).

on the available information¹²⁵ – but erroneous cascades nonetheless developed. Here is an actual example of a cascade producing an inaccurate outcome (the urn used was B)¹²⁶:

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. “[I]nitial misrepresentative signals start a chain of incorrect decisions that is not broken by more representative signals received later.”¹²⁷ 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 that 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 that 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, to be ignorant or indifferent to environmental protection. If Albert and Barbara seem to agree that global warming is a serious problem, Cynthia might not contradict them publicly and 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 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 that they are wrong. The apparent views of Albert, Barbara, and Cynthia carry information; that apparent view might be right. But even if David thinks that they are wrong, and has information supporting that conclusion, he might be most reluctant to take them on publicly. In the actual world of group decisions, people are of course uncertain whether publicly expressed statements are a product of independent knowledge, participation in an

¹²⁵ Thus 72% of subjects followed Bayes’ rule in the Anderson/Holt experiment, and 64% 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* 291, 304 (Patrick Cohendet et al. eds. 1996).

¹²⁶ See 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* 291 (Patrick Cohendet et al. eds.) (New York: Springer Verlag, 1996).

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 a clever variation on the urn experiment mentioned above.¹²⁸ In this experiment, people were paid twenty-five cents for a correct decision, but seventy-five cents 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 twenty-five cents; if their decision failed to match the group's decision, they lost seventy-five cents.

In this experiment, cascades appeared almost all of the time. No fewer than 96.7% of rounds resulted in cascades, and 35.3% 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% of people matched the first announcement. Consider, as a dramatic illustration, this period of the experiment¹²⁹ (the actual urn for this period was B):

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 that they actually have.

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 in a more extreme position in line with their tendencies before deliberation began.¹³⁰ Group polarization is the typical pattern with deliberating groups. It has been found in hundreds of studies involving over a dozen countries, including the United States, France, Afghanistan, and Germany.¹³¹ 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

¹²⁷ Anderson and Holt, *supra* note, at 859.

¹²⁸ See Hung and Plott, *supra* note, at 1515-1517.

¹²⁹ *Id.* at 1516.

¹³⁰ See Roger Brown, *Social Psychology: The Second Edition* 203-226 (New York: The Free Press, 1985).

¹³¹ See *id.* at 204.

evidence of the latter phenomenon among citizens of France.¹³²

Group polarization occurs for issues of fact as well as issues of value, though it is easiest to demonstrate for 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. But if the question, asked on a bounded scale, is whether a terrorist attack is likely to occur in the United States in the next year, the answers, from a deliberating group, will reveal group polarization, as people move toward more extreme points on the scale, depending on their initial median point. 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.¹³³ Juries polarize as well.¹³⁴

Why does group polarization occur? There are three reasons.¹³⁵ 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 the arguments 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, to hold onto their preserved self-presentation. They shift accordingly.¹³⁶ 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.¹³⁷ In a wide

¹³² *Id.* at 224.

¹³³ See Sunstein, Schkade, and Ellman, *supra* note.

¹³⁴ See David Schkade et al., *Deliberating About Dollars: The Severity Shift*, 100 *Colum. L. Rev.* 101 (2000).

¹³⁵ See Brown, *supra* note, at 200-45.

¹³⁶ *Id.* 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 people, and that when minorities have influence, it is because they produce genuine attitudinal change. See Baron et al., *supra* note, at 82.

¹³⁷ See Robert Baron et al., *Social Corroboration and Opinion Extremity*, 32 *J Experimental Soc. Psych.* 537 (1996).

variety of experimental contexts, people's opinions have been shown to become more extreme simply because their view has been corroborated, and because they have been more confident after learning of the shared views of others.¹³⁸

Does group polarization led to accurate or inaccurate answers? Do deliberating groups err when they polarize? No general answer would make sense. Everything depends on the relationship between the correct answer and the group's predeliberation tendency. But as a result of the relevant influences, some people will fail to disclose what they know. When polarization is involved, deliberative processes might well fail to move people in the right directions. When individuals are leaning in a direction that is mistaken, the mistake will be amplified by group deliberation. We have already encountered an example: When most people are prone to make conjunction errors, group processes lead to more errors rather than fewer.¹³⁹ This is polarization in action, and it produces major blunders.

Deliberative Success

Thus far I have emphasized several reasons why deliberation often fails to improve on the judgments of statistical groups, and indeed might make those judgments even worse. But there is some intriguing countervailing evidence.

When one or more people in a group are confident that they know the right answer to a factual question, the group might be expected to shift in the direction of accuracy.¹⁴⁰ Suppose that the question is how many people were on the earth in 1940, or the number of Supreme Court decisions invalidating acts of Congress, or the distance between Paris and London. Suppose too that one or a few people know the right answer. If so, there is a good chance that the group will not polarize, but instead converge on that answer. When this is so, the reason is simple: The person who is confident that he knows the answer will speak with assurance and authority, and she is likely to be convincing for that very reason. An early study finds that those with correct answers are usually more confident, and hence confidence "is associated with correctness for both individual and group performance."¹⁴¹ Consider in this light the finding that pairs tend to do better than individuals on a test involving general vocabulary knowledge; those pairs with at

¹³⁸ Baron et al., *supra* note 137.

¹³⁹ See note *supra*.

¹⁴⁰ See James Fishkin and Robert Luskin, *Bringing Deliberation to the Democratic Dialogue*, in *The Poll With A Human Face* 3, 29-31 (Maxwell McCombs and Amy Reynolds eds. 1999).

least one high-ability member generally performed at the same level as their more competent member.¹⁴²

Some evidence suggests that while deliberating groups often fail to spread information, they are less likely to neglect unshared information if they believe that there is a demonstrably correct answer to the question that they are trying to answer.¹⁴³ Asked to solve a murder mystery, a deliberating group did far better when its members were told that they had sufficient clues to “determine” the identify of the guilty suspect than when they were told to decide which suspect was “most likely to have committed the crime.”¹⁴⁴ Hence “adequate consideration of unshared, critical information during group discussion” appears to be affected by “how members construe their decision-making task,”¹⁴⁵ so that those who believe that they are solving a problem with a correct solution are more likely to explore shared information than those who think that there are reaching a consensus. It follows that “discussions may be more data driven and less consensus driven when members believe that a demonstrably correct decision exists.”¹⁴⁶ Even here, however, the member finding the right solution requires some initial support in the group; otherwise the group will frequently fail.¹⁴⁷

An impressive study finds that groups performed exceedingly well, far better than individual members, in two complex tasks that had demonstrably correct solutions.¹⁴⁸ The first involved a statistical problem, requiring subjects to guess the composition of an urn containing blue balls and red balls. The second involved a problem in monetary policy, asking participants to manipulate the interest rate to steer the economy in good directions. People were asked to perform as individuals and in groups. The basic results for the two experiments were similar. Groups significantly outperformed individuals. On a scale of 1-100, the average group score in the urn test was 86.8, as opposed to 83.7 for individuals – a highly significant difference

¹⁴¹ See Hastie, *Experimental Evidence on Group Accuracy*, *supra* note, at 148.

¹⁴² *Id.*

¹⁴³ See Laughlin, *Social Combination Processes of Cooperative Problem-Solving groups on Verbal Intellectual Tasks*, in M. Fishbein, 1 *Progress in Social Psychology* 127 (1980); Laughlin and Ellis, *Demonstrability and Social Combination Processes on Mathematical Intellectual Tasks*, 22 *J Exp Social Psych* 177 (1986).

¹⁴⁴ See Garold Strasser and Dennis Stewaty, *Discovery of Hidden Progiles by Decision-Making Groups: Solving a Problem Versus Making a Judgment*, 63 *J Personality and Social Psych* 426 (1992).

¹⁴⁵ *Id* at 432.

¹⁴⁶ *Id* at 433.

¹⁴⁷ See Robert J. MacCoun, *Comparing Micro and Macro Rationality*, in *Judgments, Decisions, and Public Policy* (M.V. Gowda and Jeffrey Fox eds. 2002).

¹⁴⁸ Alan Blinder and John Morgan, *Are Two Heads Better Than One? An Experimental Analysis of Group Vs. Individual Decisionmaking*, NBER Working Paper 7909 (2000).

statistically. For the monetary policy problem, the difference was essentially identical. Interestingly, groups did not, on balance, take longer to make decision. In terms of both accuracy and time, there were no differences between group decisions made with a unanimity requirement and group decisions made by majority rule.

How can these results be explained? An obvious possibility is that group processes play a little role and that the group's discussion is simply the average of individual judgments. On this view, the judgments of these deliberating groups simply were statistical judgments. But the evidence is inconsistent with this hypothesis; groups did far better than their average member. Even more remarkably, the performance of the median player did not explain the performance of the group. An alternative hypothesis is that each group contained one or more strong analysts, who were able to move the group in the right direction. But in the experiments, there is little support for this hypothesis. "In the end, we are left to conclude that neither the average player, nor the median player, nor the best player determine the decisions of the group." It seems that in these experiments, the better decisions by groups resulted from the fact that the best points and arguments turned out to spread among the various individual players. Here we find some basis for the claim that under appropriate conditions, groups can do much better than individuals.¹⁴⁹ The relevant conditions appear to include highly competent group members attempting to solve statistical problems that all members knew to have demonstrably correct answers.

4. Remedies and Reforms

How might group performance be improved? How can groups counteract the problems I have emphasized? If mistakes come from informational and reputational pressure, then the solution is to take steps to elicit the information that people actually have. The most difficult problem is the propagation of error. If group members use the availability heuristic, or if they fall prey to optimistic bias, blunders will result unless they are corrected by one or more group members. Even here, the best solution is to attempt to ensure that group members disclose what they know.

But for those who seek to diminish the effects of informational pressure and social influences, there is a cautionary note. We can imagine groups that actually benefit from both of

¹⁴⁹ As I have suggested, the overall evidence on this point is mixed. See Kerr et al., *supra* note.

these, and hence from cascades and polarization.¹⁵⁰ Sometimes it is good for people to silence themselves; sometimes their contributions would be unhelpful, because what they believe that they know is false.¹⁵¹ If some group members have a bad idea about how to stabilize the economy, litigate a case, or reduce the threat of terrorism, informational pressure and social influences might make them defer to those who know much better. As a result, the group will do better rather than worse.

We have seen that polarization might lead people in the right direction; the question is whether a more extreme version of members' antecedent tendency is correct, and that question must be answered on its merits. The process of polarization does not provide that answer. Or consider a cascade in which the early movers actually know the truth, and those who follow them are ignoring private information that they believe to be true but that would, on reflection, turn out to be erroneous or misleading. If so, the followers are not only rational in disregarding what they know; they also lead the group in a better direction because they do not give it bad signals. Those who participate in cascades are acting rationally, but the more important point is that if those who start cascades are correct, both individuals and groups are better off as a result. The only problem – and it is a serious one -- is that many cascade participants will fail to disclose accurate information, and for that reason the group will suffer, as demonstrated by the experiments discussed above.

The overriding question is how to alter people's incentives, so as to ensure that people will say what they actually know. Many possibilities might be imagined here.

Predeliberation Anonymity: Secret Ballots and the Delphi Method

People might be asked to register their opinions anonymously, either in advance of deliberation or after it has occurred. The secret ballot can be understood as an effort to insulate people from reputational pressures and to permit them to say what they believe.¹⁵² Many institutions should consider more use of the secret ballot simply to elicit more information.

Consider the Delphi Technique, which has several key features.¹⁵³ First, it ensures the anonymity of all members through a self-administered questionnaire. The purpose of anonymity

¹⁵⁰ Hidden profiles, of course, are never desirable, assuming the truth of the information that is hidden.

¹⁵¹ Of course falsity can sometimes contribute to truth; but it frequently does not.

¹⁵² See Timur Kuran, *Private Truths, Public Lies* (1998).

¹⁵³ See Gene Rowe and George Wright, *Experts Opinions in Forecasting: The Role of the Delphi Technique*, in

is precisely “to diminish the effects of social pressures, as from dominant or dogmatic individuals, or from a majority.”¹⁵⁴ Second, it is iterated, and there is a system for controlled feedback on the judgments of others. Members make individual estimates; all members are informed of the views of other members; and there are additional rounds of estimates, allowing feedback until there is a desired level of convergence. Third, group members are permitted to communicate, but sometimes only their ultimate conclusions (generally in the form of summary statistics involving quartiles or ranges); and typically the conclusions, given anonymously, are provided to others by a facilitator or monitor team, often in the form of a simple summary such as a mean or median value of the group response. Thus “the feedback comprises the opinions and judgments of all group members and not just the most vocal.”¹⁵⁵ (Note here that the Delphi Method is most successful when group members are provided not only with the mean or median estimate, but also with reasons given by group members for their views.¹⁵⁶ An account of reasons is most likely to move people in the correct directions.¹⁵⁷) Fourth, and finally, the judgments of group members are subject to a statistical aggregation.

The Delphi Method provides a sharp contrast with efforts to obtain the judgments of statistical groups and also with interacting groups containing open deliberation. And in several contexts, the Delphi Method has provided more accuracy than open discussion.¹⁵⁸ For general almanac questions, the Delphi Method was found to produce better answers than individual estimates, though open discussion did still better, apparently because it served to correct errors.¹⁵⁹ A natural alternative to the Delphi Method would be a system in which ultimate judgments were stated anonymously, but only after deliberation. Anonymity would insulate group members from reputational pressure, and to that extent could reduce the problem of self-silencing. But it would do little to reduce informational pressure.

Roles, Experts, and Forewarning

Imagine a deliberating group consisting of people with specific roles, appreciated and known by all group members. One person might be understood to have medical expertise; a second might be a lawyer; a third might know about public relations; a fourth might be a

Principles of Forecasting 125 (J. Scott Armstrong ed. 2001).

¹⁵⁴ Id. at 126.

¹⁵⁵ Id. at 126.

¹⁵⁶ Id. at 129.

¹⁵⁷ Id. at 129-30.

¹⁵⁸ See id; Hastie, *supra* note.

statistician. In such a group, it might be hypothesized that sensible information aggregation would be far more likely, simply because each member knows that each other has something particular to contribute. Hidden profiles should be less likely to remain hidden if there is a strict division of labor, in which each person is knowledgeable, and known to be knowledgeable, about something in particular.¹⁶⁰

Several experiments support the hypothesis.¹⁶¹ In one such experiment, each member of a three-person group was given a good deal of information about one of three candidates for office.¹⁶² In half of these groups, the “expertise” of each member was publicly identified to all before discussion began; in half of them, there was no such public identification of experts. The bias in favor of shared information was substantially reduced in those groups in which experts were publicly identified as such.¹⁶³ The reduction of the bias was significantly smaller when there was no public identification of experts and when each group member was simply told, by the experimenter, that he or she was an expert on a particular candidate.¹⁶⁴ The lesson is clear: If a group seeks to obtain the information that its members hold, it would make sense to inform all group members, before deliberation begins, that different members have different, and relevant, information to contribute. Unfortunately, however, the effect of role assignment, in reducing hidden profiles, is not huge.¹⁶⁵

Devil’s Advocates

If hidden profiles and self-silencing are the source of group failure, then an obvious response is to ask some group members to act as “devil’s advocates,” urging a position that is contrary to the group’s inclination.¹⁶⁶ This was a central suggestion of both the Senate Committee reporting on intelligence failures in connection with Iraq and of the review board

¹⁵⁹ Id.

¹⁶⁰ See Stasser, *The Uncertain Role of Unshared Information in Collective Choice*, *supra*, at 56-57.

¹⁶¹ See Garold Stasser et al., *Pooling Unshared Information: The Benefits of Knowing How Access to Information Is Distributed among Group Members*, 82 *Org Behavior and Human Decision Processes* 102 (2000); Garold Stasser et al., *Expert Roles and Information Exchange During Discussion: The Importance of Knowing Who Knows What*, 31 *J Experimental Social Psych* 244 (1995).

¹⁶² See Stasser et al., *Pooling Unshared Information*, *supra* note.

¹⁶³ Id. at 110-111.

¹⁶⁴ In this condition, there was no effect on collective recall of candidate characteristics, but this information did make it more likely that unshared information would be retained in a written protocol after it had been mentioned during discussion. Id. at 109-110.

¹⁶⁵ Stasser and Titus, *supra* note, at 310.

¹⁶⁶ See Janis, *supra* note, at 267.

that investigated large blunders at NASA.¹⁶⁷

Those assuming the role of devil's advocates will not occur the reputational pressure that comes from rejecting the dominant position within the group; they have been requested to do precisely that. And because they are asked to take a contrary position, they are freed from the informational influences that can lead to self-silencing. Hidden profiles are less likely to remain hidden if one or more group members are told to disclose the information they have, even if that information runs contrary to the apparent tendency within the group. In at least one well-known case, this approach appeared to work. "During the Cuban missile crisis, President Kennedy gave his brother, the Attorney General, the unambiguous mission of playing devil's advocate, with seemingly excellent results in breaking up a premature consensus."¹⁶⁸

Unfortunately, research on devil's advocacy in small groups does not provide conclusive evidence of the effectiveness of devil's advocacy in real-world settings.¹⁶⁹ To be sure, many experimenters have found that protection of genuine dissenting views can enhance group performance.¹⁷⁰ But a formal requirement of devil's advocacy enhances group performance far less than does the articulation of genuine dissent. When an advocate's challenges to a group consensus are insincere, members discount his arguments accordingly. At best, he merely facilitates a "multisided examination of the problems at hand."¹⁷¹ Because devil's advocates have no incentive to sway the group's members to their side, they accomplish their task if they allow the consensus view to refute the unpopular dissenting arguments. Unlike a genuine dissenter, the devil's advocate has little to gain by zealously challenging the dominant view and as a result tends not to persist in challenging the consensus.¹⁷² In any case the perceived sincerity of a dissenter is an important factor in determining minority influence.¹⁷³ An insincere devil's advocate is unlikely to provide much help.

Enlisting High-Status Contrarians – and Leadership

Some people are more likely to silence themselves than others. For example, group

¹⁶⁷ See notes supra.

¹⁶⁸ Janis, supra note, at 268.

¹⁶⁹ Gary Katzenstein, The Debate on Structured Debate: Toward a Unified Theory, 66(3) *Org. Beh. & Hum. Decision Processes* 316, 317-18 (1996).

¹⁷⁰ George Alexander & Eric Stern, Harnessing Conflict in Foreign Policy Making: From Devil's to Multiple Advocacy, 32 *Presidential Studies Quarterly* 484 (2002).

¹⁷¹ *Id.*

¹⁷² *Id.*

members are less likely to conform if they have high social status or are extremely confident about their own views.¹⁷⁴ In a complementary finding, members of low status groups – less educated people, African-Americans, sometimes women – have been shown to carry less influence within deliberating groups than their higher-status peers.¹⁷⁵ Creative groups would do well to exploit these findings.

For example, the problem of unshared information is reduced when that information is held by a leader within a group; not surprisingly, the leader's words count, and people listen to what leaders have to say.¹⁷⁶ In a leading experiment, a medical team consisting of a resident physician, an intern, and a third-year medical student showed a tendency to repeat unshared items emphasized by the resident.¹⁷⁷ More generally, those experienced in the task at hand are more likely to mention and to repeat unshared information.¹⁷⁸ One reason for these effects is that those with higher status or competence are less subject to the reputational pressures that come from emphasizing unshared information.¹⁷⁹ Another reason is that leaders and experts are more likely to think that their own information is accurate and worth disclosing to the group, notwithstanding the fact that the information held by other group members cuts in the other direction.

The simplest lesson is that leaders and high-status members can do groups a large service by asserting a contrary view, at least for purposes of argument.¹⁸⁰ In a similar vein, group leaders should be reluctant to state a firm view at the outset and should thus allow space for more information to emerge.

Restructured Incentives, 1: Money

How would groups perform if individuals knew that they would be rewarded, not if their own answer was correct, but if the majority of the group was correct? We might speculate that in a situation of this kind, hidden profiles and cascades would be dramatically reduced. The reason

¹⁷³ Moscovici at 359-65.

¹⁷⁴ See Robert Baron et al., *Group Process, Group Decision, Group Action*, supra note, at 66.

¹⁷⁵ See Caryn Christenson and Ann Abbott, *Team Medical Decision Making*, in *Decision Making in Health Care* (Gretchen Chapman and Frank Sonnenberg eds.) (New York: Cambridge University Press, 2000), at 267, 273-76.

¹⁷⁶ Stasser, *Unshared Information*, supra note, at 65.

¹⁷⁷ *Id.*

¹⁷⁸ *Id.*

¹⁷⁹ *Id.*

¹⁸⁰ Cf. Janis, supra note, at 262-63 (emphasizing the need for leaders to be willing to accept criticism of his or her own judgments).

is that when people are rewarded when their group is right, they are far more likely to reveal, to that group, what they actually know.

For supportive evidence, consider an intriguing variation on the urn experiment, where subjects were paid \$2 for a correct group decision and penalized \$2 for an incorrect group decision, with the group decision determined by majority rule.¹⁸¹ People were neither rewarded nor punished for a correct individual decision. The result was that in 92% of cases, people's announcement matched their private draw. And because people revealed their private signals, the system of majority rule produced a huge increase in fully informed decisions -- that is, the outcomes that someone would reach if he were somehow able to see all private information held by group members. As an example, consider this period from the majority rule experiment¹⁸² (the actual urn was A):

What is the explanation for this significantly reduced level of cascades in a system of majority rule? The answer lies in the fact that the individual knows that he has nothing to gain from a correct individual decision and everything to gain from a correct group decision. As a result, it is in the individual's interest to say exactly what he sees, because it is the accurate announcement, from each person, that is most likely to promote an accurate group decision. A simple way to understand this point is to assume that a group has a large number of members and that each member makes an announcement that matches his private draw. As a statistical matter, it is overwhelmingly likely that the majority's position will be correct.

The sophisticated participants in this experiment, from the California Institute of Technology, saw the point; perhaps other participants would not, and hence the result might not be generalizable. But we could imagine institutional design that would increase the likelihood that people would disclose what they know, so as to reduce the informational and reputational pressures imposed by the expressed views of others. Institutional leaders might reward people for disclosing information that is held by only a few; they might make clear, in advance, that departures from the prevailing wisdom are welcome. Approaches of this kind would increase the likelihood that hidden profiles would come to the attention of the group as a whole.

Restructured Incentives, 2: Norms

¹⁸¹ Plott and Huang at 1511.

¹⁸² Id. at 1515.

Self-silencing is partly a product of social norms – of a sense that people will be punished, rather than rewarded, for disclosing information that departs from the group’s inclination. It should be easy to see that groups can aggravate or eliminate this effect. If consensus is prized, and known to be prized, then self-silencing will be more likely. If the group is known to welcome new and competing information, then the reward structure will be fundamentally different. Evidence for this claim comes from experiments that “primed” people by asking them to engage in a prior task that involved either “getting along” or “critical thinking.” Primed by a task that called for critical thinking, people were far more likely to disclose what they know, and there was a quite substantial effect on hidden profiles.¹⁸³ The general lesson is that if norms favor disclosure of privately held information, then self-silencing will be significantly reduced; deliberation is likely to benefit as a result.

5. Information Markets

Deliberation is one way to aggregate the information held by group members; another way is to rely on the price signal, which has a similar aggregative function.¹⁸⁴ And if an emphasis is placed on the information-aggregating properties of markets, it would seem plain that if we are attempting to improve on the answer produced by statistical means and deliberating groups, we might consider an increasingly popular possibility: *Create a market*.¹⁸⁵ Information markets have proved remarkably success at forecasting future events; they seem to do far better, in many domains, than deliberating groups.

A central advantage of information markets is that they impose the right incentives for people to disclose the information that they hold. Because investments are generally not disclosed to the public, investors need not fear reputational sanctions if, for example, they have

¹⁸³ Stasser and Titus, *supra* note, at 143.

¹⁸⁴ See F.A. Hayek, *The Use of Knowledge in Society*, 35 *Am Econ Rev* 519 (1945). For discussion of information markets (sometimes called prediction markets), see Joyce Berg et al., *Results from a Dozen Years of Election Futures Markets Research*, in *Handbook of Experimental Economic Results* (Charles Plott and Vernon Smith eds 2003); R. Forsythe et al., *Anatomy of an Experimental Political Stock Market*, 82 *Am Ec Rev* 1142 (1992); Joyce Berg et al., *What Makes Markets Predict Well? Evidence from the Iowa Electronic Markets*, in *Understanding Strategic Interaction* (W. Albers et al. eds. 1997); R. Forsythe et al., *Wishes, Expectations, and Actions: Price Formation in Election Stock Markets*, 39 *J Ec Behavior and Org* 83 (1999).

¹⁸⁵ For valuable overviews, see Justin Wolfers and Eric Zitzewitz, *Prediction Markets*, 18 *J Econ Persp* 107 (2004); Michael Abramowicz, *Information Markets, Administrative Decisionmaking, and Predictive Cost-Benefit Analysis*, 71 *U Chi L Rev* 933 (2004); Saul Levmore, *Simply Efficient Markets and the Role of Regulation*, 28 *J Corp Law* 589 (2003).

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. The use of that information will be reflected in the price signal. 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. For small groups, of course, information markets are likely to be too "thin" to be useful; a certain number of investors is required to get a market off the ground.¹⁸⁶ In many contexts, however, private and public organizations might use markets as a complement to or even a substitute for deliberation.

Practice and Evidence

An abandoned initiative

In many imaginable markets, people might make claims about facts, or predictions about the future, and they might stand to gain or lose from their predictions. In the summer of 2003, analysts at the Department of Defense built directly on this idea.¹⁸⁷ To predict important events in the world, including terrorist attacks, they sought to create a kind of market in which ordinary people could actually place bets. The proposed Policy Analysis Market would have allowed people to invest in their predictions about such matters as the growth of the Egyptian economy, the death of Yasser Arafat, the military withdrawal of the United States from specified nations, and the likelihood of terrorist attacks in the United States. Investors would have won or lost money on the basis of the accuracy of their predictions.

Predictably, the Policy Analysis Market produced a storm of criticism. Ridiculed as "offensive" and "useless," the proposal was abandoned. Senator Tom Daschle called the market "a plan to trade in death" and boldly claimed that the plan was "the most irresponsible, outrageous and poorly thought-out of anything that I have heard the administration propose to date."¹⁸⁸ Senator Byron Dorgan argued that "it is morally bankrupt for a government agency to make a profitable game out of the deaths of American troops, heads of state, and nuclear missile

¹⁸⁶ But see Levmore, *supra* note (showing considerable success even within quite thin markets).

¹⁸⁷ See Justin Wolfers and Eric Zitzewitz, Prediction Markets, 18 J Econ Persp 107 (2004).

¹⁸⁸ R. Bailey, Betting on Terror, *Reason Online*, <http://www.reason.com/rb/rb073003.shtml>

attacks.”¹⁸⁹ A private Policy Analysis Market, specializing in the Middle East, was promised in 2003, but it did not go forward.¹⁹⁰

Amid the war on terrorism, why was the Defense Department so interested in the Policy Analysis Market? The answer is simple: it wanted to have some help in predicting geopolitical events, including those that would endanger American interests, and it believed that a market would provide that help. It speculated that if a large number of people could be given an incentive to aggregate their private information, in the way that the Policy Analysis Market would do, government officials would learn a great deal. Apparently it believed that such a market would provide an important supplement to deliberative processes within government and without.¹⁹¹

Iowa Electronic Markets

Does this idea seem fanciful? Since 1988, the University of Iowa has run the Iowa Electronic Markets (IEM), which allow people to bet on the outcome of presidential elections. Originally the IEM allowed people to trade only in the expected fraction of the popular vote to be obtained by presidential candidates.¹⁹² Securities were offered that would pay \$2.50 multiplied by the specified candidate’s share of the vote. If, for example, George H.W. Bush received 50% of the vote, then the shareholder would receive \$1.25. Shares could be bought and sold until the day before the election. Since their opening, the IEM have expanded from these modest roots. In the recent past, traders have been able to bet on the market capitalization that Google will achieve in its initial public offering, the price of Microsoft stock at a future date, and Federal Reserve monetary policy, in addition to betting on American elections.¹⁹³

For presidential elections—still the most popular markets that IEM operates—traders can now choose from two types of markets.¹⁹⁴ In a “winner-take-all” market, traders win \$1 for each future in the winning candidate that they own and nothing for shares of the losing candidate. In a “vote-share” market, traders in “candidate futures” win \$1 multiplied by the proportion of the

¹⁸⁹ B. Dorgan, The Pentagon’s Ill-Conceived Market, *Wash. Post*, Aug. 7, 2003 at A20.

¹⁹⁰ For a replicate of the site, see <http://www.ratical.org/ratville/CAH/linkscopy/PAM/>

¹⁹¹ See Wolfers and Zizewitz, *supra* note.

¹⁹² See Joyce Berg et al., Accuracy and Forecast Standard Error of Prediction Markets (July 2003 working paper).

¹⁹³ See <http://www.biz.uiowa.edu/iem/markets/>

¹⁹⁴ See *id.*

popular vote that the candidate received.¹⁹⁵ Thus, in a winner-take-all market, a “Dukakis future” was worth nothing after the election, while in a vote-share market, each Dukakis future paid \$0.456. In a winner-take-all market, the market price reflects traders’ perception of the likelihood that each candidate will win the election. Similarly, observers can use the prices in a vote-share market much as they might use a poll. These prices are the market’s estimate of each candidate’s likely share of the vote when the election occurs. In each case, the market price reflects the aggregate information held by participants.

The IEM operate much like an ordinary stock market. To enter, each participant must purchase “unit portfolios” consisting of one future in each candidate for each dollar that the trader puts into the market.¹⁹⁶ Once she has bought enough of these “unit portfolios,” she can unbundle the contracts and trade individual shares. All trading is fully computerized and traders must reach the markets through the Internet.¹⁹⁷ Unlike most stock exchanges, the IEM does not allow speculators to sell futures short. Nevertheless, as in a typical stock market, traders can issue bids and asks (limit orders) or accept outstanding offers (market orders). While most traders merely accept market orders rather than choosing their own prices, a small group of “marginal traders” trade frequently and post limit orders.¹⁹⁸ It is these traders who have the greatest effect on prices, as we shall see.

As a predictor, the Iowa Electronic Markets have produced extraordinarily accurate judgments. Most of the time, they have done better than professional polling organizations.¹⁹⁹ In the week before the last four elections, the predictions in the Iowa market have shown an average absolute error of just 1.5 percentage points, a significant improvement over the 2.1 percentage point error in the final Gallop polls.²⁰⁰ The IEM have proved accurate not only on election eve but only in long forecasting horizons, both in absolute terms and also when compared to alternative forecasting systems.²⁰¹ Nor are such markets limited to the United States. In other nations, universities are operating similar markets; examples include the University of British

¹⁹⁵ Robert Forsythe, Thomas Rietz, & Thomas Ross, Wishes, Expectations, and Actions: A Survey on Price Formation in Election Stock Markets, 39 *J. Econ. Behav. & Org.* 83, 85 (1999).

¹⁹⁶ *Id.* at 86.

¹⁹⁷ *Id.*

¹⁹⁸ *Id.* at 99-100.

¹⁹⁹ See Wolfers and Zitzewitz, *supra* note, at 112.

²⁰⁰ *Id.*

²⁰¹ See Joyce Berg et al., Accuracy and Forecast Standard Error of Prediction Markets (July 2003 working paper).

Columbia Election stock market, involving Canada,²⁰² and Vienna University of Technology, operating the Austrian Electronic Market.²⁰³ Although the relevant districts are quite small, Australian bookmakers have shown a high degree of accuracy in predicting district-level races.²⁰⁴

Other information markets

Outside of the political context, consider the Hollywood Stock Exchange, in which people predict Oscar nominees and winners (as well as opening weekend box office successes). For the Hollywood Stock Exchange, the level of accuracy has been impressive. “HSX offers good predictions of a film's gross receipts before release and, relatively speaking, even better predictions after opening weekend - when a large number of traders have some information in the form of (or at least the possibility of) observing the finished film on screen, along with audience reactions. Apparently, studios have begun relying on these estimates to structure the distribution of their films.”²⁰⁵ The market has proved at least equal to expert panels in predicting Oscar winners, with (for example) correct predictions of thirty-five of forty Oscar nominees in 2002.²⁰⁶

The futures market for oranges does a better job predicting weather in Florida than the National Weather Service.²⁰⁷ A large prediction market, producing a typical event turnover in the hundreds of millions of dollars and run by the Deutsche Bank and Goldman Sacks, involves the likelihood that economic data released later in the week will show specific values²⁰⁸; the market performs about as well as the consensus forecasts of a survey of about fifty professional forecasters.²⁰⁹ Companies have started to use internal prediction markets to answer relevant questions, including likely sales in specific periods.²¹⁰ The level of accuracy here is also high – far better, in fact, than what would emerge from statistical means or deliberation, where

²⁰² <http://esm.ubc.ca/>

²⁰³ <http://ebweb.tuwien.ac.at/apsm/>

²⁰⁴ Justin Wolfers and Andrew Leigh, Three Tools for Forecasting Federal Elections: Lessons from 2001, 37 *Australian Journal of Political Science* 223 (2002).

²⁰⁵ Levnore, *supra* note.

²⁰⁶ David Pennock et al., The Real Power of Artificial Markets, 291 *Science* 987 (2001).

²⁰⁷ R. Roll, Orange Juice and Weather, 74 *Am. Econ. Rev.* 861 (1984).

²⁰⁸ See www.economicderivatives.com

²⁰⁹ See Wolfers and Zitzewitz, *supra* note, at 114.

²¹⁰ See Charles Plott, Markets as Information Gathering Tools, 67 *Southern Economic J.* 1 (2000).

excessive optimism can cause serious problems.²¹¹

For example, Hewlett Packard (HP) and the California Institute of Technology initiated a project to study experimental markets as an information aggregation mechanism involving product sales.²¹² The experimenters chose twelve people who worked in different parts of HP's business operation. Because of its small size, the market was a very "thin" one, meaning that there were few participants and that the market was far less liquid than the much "thicker" Iowa Electronic Markets. Participants were chosen with the thought that each could contribute information from his department in buying and selling the relevant futures, which were tied to sales and bonuses for executives (which, in turn, are closely tied to profits). The markets were organized so that securities existed for intervals of sales. For example, one security would pay off if sales were between one and ten printers; another would pay off if sales were between 10 and 20. In most of the experiments, the possible range of sales was divided into ten intervals of equal size. On the basis the prices of each security, the experimenters could guess how many units HP would sell that month.

The results showed that the markets' predictions were a considerable improvement over HP's official forecast. In six of the eight markets for which official forecasts were available, the market prediction was closer to the actual outcome than the official forecast²¹³ -- and this was despite "anecdotal evidence" that the markets' activities were included as inputs in generating the official forecast.²¹⁴

In fact prediction markets are springing up all over the Internet, allowing people to make bets on the likely outcomes of sports, entertainment, finance, and political events. In fact we can find actual or proposed prediction markets about any number of questions: Will gas prices reach \$3 per gallon? Will cellular life be found on Mars? Will Osama Bin Laden be captured by a certain date? Will small pox return to the United States? Will there be a sequel to Master and Commander? Will the Federal Communications Commission be abolished? Consider the following list:

²¹¹ Id.

²¹² K. Chen and C. Plott, Information Aggregation Mechanisms: Concept, Design, and Implementation for a Sales Forecasting Problem; working paper; available online at http://www.hp1.hp.com/personal/KayYut_Chen/paper/ms020408.pdf

²¹³ Id. at 12.

²¹⁴ Id. at 5.

- Hollywood Stock Exchange— <http://www.hsx.com>
- Austrian Electronic Markets--<http://ebweb.tuwien.ac.at/apsm/>
- University of British Columbia Election Stock Market--<http://esm.ubc.ca/>
- Iowa Electronic Markets—<http://www.biz.uiowa.edu/iem/>
- Foresight Exchange—<http://www.ideosphere.com/fx/>
- Tradesports—<http://www.tradesports.com>
- Centrebet-- <http://www.centrebet.com/>
- News Futures--<http://us.newsutures.com/home/home.html>
- Probability Sports—<http://www.probabilitysports.com>
- Economic Derivatives—<http://www.economicderivatives.com>
- Wahlstreet—German political futures market;
<http://tagesspiegel.wahlstreet.de/share/home/home.html>

Aggregating information through markets

All in all, prediction markets have been spectacularly successful in terms of the aggregate accuracy of the resulting “prices.” Why is this? Note that they do not rely on the median or average judgment of a randomly selected group of people. They are genuine markets. Those who participate are self-selected. They must believe that they have relevant information; it is costly for them to “vote,” and they ought not to be expected to do so unless they have something to gain.²¹⁵ In addition, votes are not weighted equally. If people want to invest a few dollars, they are permitted to do so, but they can invest a great deal more if they are confident of their answer.²¹⁶ Intensity of conviction is captured in prices.

There is a further point. People are permitted to buy and sell shares on a continuing basis. “Unlike polls or expert panels in which participants are asked for their independent opinions, each trader in the market sees the net effects of the beliefs of all other traders, and the time series changes in those beliefs. This makes the market more than a static, one-time prediction but rather

²¹⁵ Note that some markets involve real rather than virtual money. Newsfutures, for example, uses virtual currency that can be redeemed for monthly prizes (such as appliances); Foresight Exchange and the Hollywood Stock Exchange use “virtual currency,” so that people do not earn real money, but instead attempt to enhance their reputation and their self-image. Note in this regard that Foreign Exchange lists publicly the “top ten investors by score,” see <http://www.ideosphere.com/fx/>

²¹⁶ Some markets, however, impose limits on permissible investments; the IEM is an example, with a ceiling of \$500.

a dynamic system that can respond instantaneously to the arrival of new information.”²¹⁷ Moreover, a correct answer is rewarded and an incorrect one is punished. Hence investors have a strong incentive to be right. In these circumstances, accurate answers can emerge even if only a small percentage of participants have good information. In the Iowa Electronic Markets, for example, it turns out that 85% of the traders do not seem to be particularly wise.²¹⁸ They hold onto their shares for a long period and then simply accept someone else’s prices. The predictions of the market are driven by the other 15% -- frequent traders who post their offers rather than accepting those made by other people. To work well, prediction markets do not require accurate judgments by anything like the majority of participants.²¹⁹ In this sense, information markets are very different from the ordinary judgments of groups.

How might institutions take advantage of information markets? We can imagine both internal and public varieties. An internal market would be limited to people within the relevant organization. As we have seen, Hewlett-Packard has used such a market to predict sales, and the Department of Defense proposed an internal Policy Analysis Market as part of its abandoned initiative on geopolitical events.²²⁰ In either case, the outcome of the market might well be more accurate than the outcome of deliberation, in which errors might arise and be propagated as a result of discussion. (For companies, optimistic bias is an obvious risk,²²¹ one that information markets should reduce.) An organization might rely on an internal market if it seeks to keep the results private or if it believes that an aggregation of information held within the organization will be sufficiently accurate. One risk of an internal market is that it might be too “thin”; another is that members of the organization might suffer from a systematic bias. Alternatively, an institution might create a public market, available to all, believing that through this route it will obtain more accurate results. In either case, an organization might use an information market instead of group deliberation, or at the very least as an input into such deliberation.

Failed Predictions?

In what circumstances might information markets fail? To answer this question, ordinary stock markets are the place to start. As in ordinary markets, a primary concern is that these

²¹⁷ Berg et al., *supra* note.

²¹⁸ See *id.*

²¹⁹ The same is of course true of ordinary markets. For a good overview, see Andrei Shleifer, *Inefficient Markets* (2000).

²²⁰ See note *supra*.

²²¹ See Daniel Kahneman and Dan Lovallo, *Timid Choices and Bold Forecasts: A Cognitive Perspective on Risk*

markets can be susceptible to manipulation by powerful speculators. The only known attempt to manipulate a prediction market occurred during the 2000 presidential election. A group of speculators mounted an attack on the Iowa Electronic Market by buying large volumes of futures in Patrick Buchanan. The value of Buchanan share did increase dramatically, but then fell almost immediately when “well-informed traders ... seized the opportunity to profit off the manipulative traders.”²²² Hence the Iowa market remained stable despite this attempted manipulation. Perhaps other, more plausible efforts at manipulation would succeed; but none has thus far.

Another concern is that some of the cognitive biases that afflict individuals will manifest themselves in prediction markets. Just as in group deliberation, investors in a market might be subject to predictable heuristics and biases. The results here are unequivocal: they are. Psychologists have found that people overestimate the likelihood that their preferred candidate will win an election – a form of optimistic bias.²²³ For example, at a certain point in the 1980 campaign, 87% of Jimmy Carter’s supporters believed that he would win, while 80% of Ronald Reagan’s supporters believed that their candidate would win.²²⁴ Obviously, at least one side had overestimated its candidate’s probability of victory at that point.

In the market context, IEM traders show the same bias. In 1988, for example, Dukakis supporters were more likely to hold futures in the Massachusetts governor’s ill-fated presidential bid than were supporters of George Bush.²²⁵ More strikingly still, Dukakis supporters were more likely to view the candidates’ debates as helpful to the Democratic candidate and accordingly bought significant additional futures in his campaign after each debate.²²⁶ Bush supporters showed the same pattern. Traders clearly exhibited the “assimilation-contrast” effect.²²⁷ People usually assimilate new information in a way that confirms their view of the world.

Despite these biases, the Iowa Electronic Markets were more accurate than most polls in predicting the outcome of the 1988 presidential election. Even three weeks before the election,

Taking, in *Choices, Values and Frames* 393 (Daniel Kahneman and Amos Tversky eds. 2000).

²²² Klarreich, Best Guess, *Science News* (Oct 18, 2003); available online at <http://www.sciencenews.org/articles/20031018/bob9.asp>.

²²³ For an overview, see Christine Jolls, Behavioral Economics Analysis of Redistributive Legal Rules, 51 *Vand. L. Rev.* 1653 (1998)

²²⁴ D. Granberg and E. Brent, When Prophecy Bends: The Preference-Expectation Link in U.S. Presidential Elections, 45 *Journal of Social and Personality Psychology* 477.

²²⁵ Forsythe, Rietz, and Ross at 94.

²²⁶ *Id.*

²²⁷ M. Sherif and C. Hovland, *Social Judgment: Assimilation and Contrast Effects in Communication and Attitude Change* (1962).

the market provided an almost-perfect guess for the candidates' shares of the vote.²²⁸ How is this result possible when many traders showed the relevant biases? The answer lies in the behavior of a small group of "marginal traders" who were far less susceptible to these biases -- the "marginal trader" hypothesis. In this view, a small group of traders who are less susceptible to these biases have a much greater effect on aggregate market behavior. In trading election futures, these traders did not show the same biases as their fellow traders and earned significant profits at the expense of their less rational colleagues.²²⁹ Thus, the biased behavior of most traders did not affect the market price because the marginal traders were prepared to take advantage of their blunders. If marginal traders are active and able to profit from the bounded rationality of other participants, then there will be no effect on the aggregate market price.²³⁰

Another bias that might be expected to affect prediction markets is the "favorite-longshot" bias often seen in horse races. In horse-racing, heavy favorites tend to give higher returns than other horses in the field, while longshots tend to offer lower than expected returns.²³¹ If the point generalizes, prediction markets might not be accurate with respect to highly unlikely events. The market should be expected to overestimate their likelihood; for example, Pat Buchanan futures would be expected to be (and might well have been) overpriced even before the attempted manipulation of the market. By contrast, a prediction market might underestimate the probability of events that are highly likely to occur.²³² But on existing prediction markets, there is little evidence of systematic errors in this vein.

"Prediction bubbles" are also easy to imagine, with investors moving in a certain direction with the belief that many other investors are doing the same. A temporary upsurge in investment in the nomination of Hillary Rodham Clinton as 2004 Democratic nominee might well have been a small bubble, with some investors thinking, not that she would in fact be the nominee, but that others would invest in that judgment, thus inflating the value of the investment. Crashes are possible as well. In any case informational influences can certainly lead individuals

²²⁸ J. Berg, F. Nelson, and T. Rietz, Accuracy and Forecast Standard Error of Prediction Markets, Working draft; available online at: <http://www.biz.uiowa.edu/faculty/trietz/papers/forecasting.pdf>

²²⁹ Forsythe, Rietz, and Ross at 100.

²³⁰ Compare the discussion in Schleifer, *supra* note.

²³¹ See Richard Thaler and William Ziemba, Anomalies: Parimutual Betting Markets: Racetracks and Lotteries, 2 J Econ Persp 161 (1988); see also C. Manski, Interpreting the Predictions of Prediction Markets, unpublished; available online at http://faculty.econ.nwu.edu/faculty/manski/prediction_markets.pdf

²³² C. Wolfers and E. Zitzewitz, Prediction Markets, preliminary draft; available online at <http://faculty-gsb.stanford.edu/wolfers/Papers/Predictionmarkets.pdf>

to make foolish investments in any market, including prediction markets.²³³

In particular contexts, the problems are worse still. Consider the problem of “terrorism futures.” It would be extremely valuable to aggregate privately held information about the risk and location of any attack. But do likely investors actually possess helpful information? Thomas Rietz, a director of the Iowa Electronic Markets, argued that terrorism and world events were fundamentally different from other contexts in which markets have successfully predicted future events.²³⁴ When betting on presidential elections, people can use their network of friends, family, and co-workers to form an opinion; but for most investors, there are no such sources of information for terrorist activity. Another skeptic worried that the market would allow the wealthy to “hedge” against the possibility of terrorist activity, while ordinary Americans would remain vulnerable to this threat.²³⁵ In this view, “terrorism futures” could operate as an insurance market that would not serve its purpose of providing information. In any event government use of the resulting information could be self-defeating, at least if the information were made public. Terrorists would know the anticipated time and location of attacks, and also know that the government was aware of this – which would make it most unlikely that the prediction would turn out to be accurate. Where the event’s occurrence is endogenous to the outcome of the information market, there is reason for skepticism about its likely performance, certainly if relevant actors have much to lose if the market turns out to be correct.²³⁶

But many policy issues, including those potentially involved in the now-defunct Policy Analysis Market, did not have this feature. Consider, for example, the question whether the Egyptian economy is likely to grow in the next year, or whether Yasser Arafat will lead the Palestinian Authority by the end of 2005. Perhaps many investors will lack a great deal of information on such questions, but it is most unlikely that the market prediction will turn out to be self-defeating. The Policy Analysis Market itself raises many questions and doubts. The broader point is that in many domains, information markets are extremely promising, and likely to outperform both statistical means and the products of group deliberation. At a minimum, such markets should be used, where feasible, as an adjunct to deliberative processes.

²³³ See Robert Schiller, *Irrational Exuberance* (2001).

²³⁴ C. Biever and D. Carrington, Pentagon cancels futures market on terror, *New Scientist Online*, <http://www.newscientist.com/news/news.jsp?id=ns99994007>

²³⁵ J. Stiglitz, Terrorism: There’s No Futures in It, *L.A. Times*, July 31, 2003 at ??; available at <http://www.commondreams.org/views03/0731-08.htm>

²³⁶ See Richard Posner, *Catastrophe: Risk and Response* (forthcoming 2004).

6. Normative Questions and Group Judgments

Deliberating groups are often asked to answer questions that are not purely factual. Issues involving morality, politics, and law require assessment of normative issues. Should cost-benefit analysis be the foundation of regulatory decisions? Should the minimum wage be increased? Should capital punishment be permitted? Can the President be impeached for lying under oath? Should *Roe v. Wade* be overruled? Should the Constitution be interpreted to require states to reconsider same-sex marriages? When, if ever, is theft morally acceptable?

When people answer such questions, informational influences and social pressures are likely to play a major role. One study demonstrates group polarization with respect to outrage: When individuals are outraged about corporate misconduct, juries are systematically more outraged than their median member.²³⁷ And in fact group discussion often produces polarization on normative issues,²³⁸ in a way that strongly suggests the presence of hidden profiles. It is on normative questions, above all, the groups end up at a more extreme point in line with their predeliberation tendencies. I have noted that in many domains, federal judges are subject to group polarization, with both Democratic and Republican appointees showing a tendency to extremism when they are sitting with like-minded others.²³⁹ It might be controversial to suggest that groups propagate individual errors, because in the normative domain, we might not be able to say, with confidence, that one or another view counts as an “error.” But if individual errors do occur, group errors will occur as well. As obvious examples, consider the persistence of slavery and racial segregation. As a less obvious example, consider framing. The framing of options affects judgments not only on factual questions but on moral ones as well, including for example the disputed issue of moral obligations to members of future generations.²⁴⁰ As noted, groups do not show less susceptibility to framing effects than individuals,²⁴¹ and hence groups will be vulnerable to framing for questions of morality and law as well as for questions of fact.

No information market could be helpful in answering normative questions, simply

²³⁷ See David Schkade et al., *Deliberating About Dollars: The Severity Shift*, 100 *Colum L Rev* 1139 (2000).

²³⁸ See Roger Brown, *Social Psychology: The Second Edition* (1985); Cass R. Sunstein, *Why Societies Need Dissent* (2003).

²³⁹ See Cass R. Sunstein et al., *Ideological Voting on Federal Courts of Appeals: A Preliminary Analysis*, 90 *Virginia Law Review* 301 (2004).

²⁴⁰ See Cass R. Sunstein, *Moral Heuristics and Moral Framing*, 88 *Minn L Rev* 1556 (2004).

because there is no way to know whether a particular investor has been shown to be correct.²⁴² And for normative questions, it might seem odd or perhaps even bizarre to rely on the judgments of statistical groups. To be sure, democratic processes might be seen as an effort to settle moral and political issues by seeking the mean view within the relevant population. But to say the least, it is controversial to claim that ordinarily moral and political questions are best answered by simply finding the mean views of a population-wide sample. (Is the morality of abortion, or capital punishment, properly assessed by asking for the average view of a group of, say, 1000 people?) Note, however, that empirical questions are often a central component of good answers to normative problems, at least on a certain view about how to think about problems of this kind, and if this is so, then the analysis of group mistakes should play to those components of group assessments of normative questions.

Consider, for example, the suggestion that the minimum wage should be increased. If minimum wage increases would significantly decrease employment, surely that is relevant to the decision whether to support such increases; and it matters too whether minimum wage increases would benefit poor people or mostly people who are not poor.²⁴³ Now these are empirical questions on which experts are almost certainly far better than deliberating groups of ordinary people. But many normative questions cannot sensibly be resolved without information about the effects of one or another answer. The argument for making cost-benefit analysis the foundation of regulatory decisions must turn, in part, on the effects of doing so. The consequences of overruling Roe v. Wade should probably matter to the decision whether that precedent should be overruled; if the result would be to increase dangerous, unlawful abortions, courts should, on a sensible view of the normative question, be reluctant to overrule it even if they would otherwise be inclined to do so.

Of course consequences may not be the central part of a normative dispute, and indeed it is difficult even to know how to evaluate or even to describe consequences without some kind of normative view. But the more general point nonetheless holds: Sometimes a certain view of the facts can bring diverse people into line on normative issues, producing a single position despite disagreements on those issues. To this extent, the analysis here applies to normative questions as

²⁴¹ See Kerr et al., *supra* note,

²⁴² It might be tempting to say that the moral views of posterity provide the relevant test, but then the bet would be on the moral views of posterity, not on what morality requires.

²⁴³ For evidence, see Daniel Shaviro, *The Minimum Wage, the Earned Income Tax Credit, and Optimal Subsidy*

well. Group judgments on such questions will be distorted by hidden profiles, the propagation of errors, and self-silencing. It is important to take steps, of the kind that I have catalogued, to reduce those distortions.

What about for purely normative issues, lacking any factual component? Here the argument on behalf of group deliberation is not fundamentally different from what it is elsewhere.²⁴⁴ Unless we are skeptics, we will agree that one point of deliberation is to ensure that normative questions are correctly answered, that is, are answered by reference to good reasons, even if we disagree about what they are. And if this is so, then there is strong reason to be concerned, for normative questions no less than empirical ones, that group judgments will be impaired by the mechanisms traced here. The structural reforms have an equivalent role in the normative domain.

7. Conclusion

Groups often contain a great deal of information, and an important task is to elicit and use the information that members actually have. 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 result, groups tend to propagate and even to amplify cognitive errors. They also emphasize shared information at the expense of unshared information; hidden profiles are a result. Cascade effects and group polarization are common.

What can be done by way of response? At the very least, it should be possible to structure deliberation so as to increase the likelihood that relevant information will emerge. Institutions might ensure anonymity and private polling before deliberation; they might permit anonymous statements of final conclusions; they might create strong incentives, economic and otherwise, to encourage people to disclose what they know. Information markets have significant advantages over deliberative processes, and they might supplement or replace those processes. Because they provide economic rewards for correct individual answers, such markets are often more accurate than the judgments of deliberating groups. To the extent feasible, such groups would often do

Policy, 64 U. Chi. L. Rev. 405, 450-51 (1997).

²⁴⁴ See Amy Gutmann and Dennis Thompson, *Democracy and Disagreement* (1999).

well to enlist information markets in arriving at their judgments.

My emphasis throughout has been on the aggregation of information and the risk that deliberating groups will err on instrumental questions and on issues of fact. But the same risks arise in the normative domain, where informational influences and social pressures also produce forms of self-silencing that are highly damaging to good deliberation. In that domain as elsewhere, incentives make all the difference; well-functioning groups take steps to ensure that on normative questions as on factual ones, people feel free to disclose what they believe to be true.

Table1: An informational cascade

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

Table 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

Table 3: No cascade

Private draw	a	a	a	a	b	a	a	a	b	
Decision	A	A	A	A	B	A	A	A	B	