


CHICAGO

JOHN M. OLIN LAW & ECONOMICS WORKING PAPER NO. 401
(2D SERIES)

PUBLIC LAW AND LEGAL THEORY WORKING PAPER NO. 215



FOUR FAILURES OF DELIBERATING GROUPS

Cass R. Sunstein and Reid Hastie

THE LAW SCHOOL
THE UNIVERSITY OF CHICAGO

April 2008

This paper can be downloaded without charge at the John M. Olin Program in Law and Economics Working Paper Series: <http://www.law.uchicago.edu/Lawecon/index.html> and at the

Public Law and Legal Theory Working Paper Series:

<http://www.law.uchicago.edu/academics/publiclaw/index.html>

and The Social Science Research Network Electronic Paper Collection:

http://ssrn.com/abstract_id=1121400

Four Failures of Deliberating Groups

Cass R. Sunstein* and Reid Hastie**

Abstract

Many groups make their decisions through some process of deliberation, usually with the belief that deliberation will improve judgments and predictions. But deliberating groups often fail, in the sense that they make judgments that are false or that fail to take advantage of the information that their members have. There are four such failures. (1) Sometimes the predeliberation errors of group members are amplified, not merely propagated, as a result of deliberation. (2) Groups may fall victim to cascade effects, as the judgments of initial speakers or actors are followed by their successors, who do not disclose what they know. Nondisclosure, on the part of those successors, may be a product of either informational or reputational cascades. (3) As a result of group polarization, groups often end up in a more extreme position in line with their predeliberation tendencies. Sometimes group polarization leads in desirable directions, but there is no assurance to this effect. (4) In deliberating groups, shared information often dominates or crowds out unshared information, ensuring that groups do not learn what their members know. All four errors can be explained by reference to informational signals, reputational pressure, or both. A disturbing result is that many deliberating groups do not improve on, and sometimes do worse than, the predeliberation judgments of their average or median member.

Many groups make their decisions through deliberation. When private companies or public officials are deciding how to proceed, they usually talk the problem through. But why, exactly, is deliberation important or even desirable?

A large part of the answer must be that deliberation will result in wiser judgments and better outcomes. But does deliberation actually have this effect? The answer is by no means clear. Group members may impose pressures on one another, leading to a consensus on falsehood rather than truth. A group of like-minded people, with similar predilections, is particularly vulnerable to this problem. The idea of “groupthink,” coined and elaborated by Irving Janis, suggests that groups may tend toward uniformity and

* Karl N. Llewellyn Distinguished Service Professor, Law School and Department of Political Science, University of Chicago.

** Robert S. Hamada Professor of Behavioral Science, University of Chicago, Graduate School of Business. In revised form, this essay will be the basis of a chapter in a forthcoming book, by Reid Hastie and Cass R. Sunstein, under the tentative title *Smart Groups* (Harvard Business School Press); we draw throughout on some of the discussion in Cass R. Sunstein, *Infotopia* (New York: Oxford University Press, 2006).

ensorship, thus failing to combine information and enlarge the range of arguments.¹ Much of the time, both private and public groups often blunder not in spite of deliberation but because of it. After deliberation, companies, labor unions, and religious organizations often make bad decisions; the same point holds for governments.

One of our principal goals here is to go beyond the idea of groupthink – to elaborate on several decades of research that explores the mechanisms that account for the failures of deliberating groups. To explain why such groups often do poorly, we investigate two sets of influences on members of deliberating groups. The first involves **informational signals**, which lead group members to fail to disclose what they know out of respect for the information publicly announced by others. The second involves **social pressures**, which lead people to silence themselves in order to avoid reputational sanctions, such as the disapproval of relevant others. As a result of these influences, groups run into four independent problems:

- ♦ They amplify rather than correct individual errors.
- ♦ They fall victim to cascade effects, as people follow the statements and actions of their predecessors.
- ♦ They tend to end up in more extreme positions in line with the predeliberation tendencies of their members (“group polarization”).
- ♦ They emphasize shared information at the expense of unshared information.

Because of these problems, groups often fail to achieve their minimal goal of aggregating the information actually held by their members. A confident, cohesive, but error-prone group is nothing to celebrate. On the contrary, it might be extremely dangerous, both to itself and to others. When groups make poor or self-destructive decisions, one of these four problems is often the explanation.

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

Prospects

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

- *Groups are equivalent to their best members.* One or more group members will often know the right answer, and other members might become convinced of this fact. For this reason, groups might perform toward or at the level of their best members. If some or many members suffer from ignorance or from a form of bias that leads to error, others might correct them. Deliberation might correct individual errors rather than propagate them, in a way that allows convergence on the judgment of the most accurate group member. Imagine, for example, that ten people are trying to make some prediction about the future, and that one of the ten is both an expert and a superb prognosticator. Perhaps the other group members will quickly see that they have an expert in their midst, and they will follow his lead. Consider “eureka” problems, in which the right answer, once announced, is clear to all. For such problems, groups should be expected to converge on the answer as announced by the member who actually knows it.
- *The whole is the sum of the parts: aggregating information.* Deliberation could help to aggregate existing information in a way that leads the group as a whole to know more than any individual member does. Suppose that the group contains no experts on the question at issue, but that relevant information is dispersed among members so that the group is potentially expert even if its members are not. Or suppose that the group contains a number of experts, but that each member is puzzled about how to solve a particular problem. Deliberation might elicit the relevant information and allow the group to make a sensible judgment. In this process, the whole is equal to the sum of the parts—the sum of the parts is what is sought—and no member has all the parts. When a group is trying to solve a crossword puzzle, something of this kind often occurs, as different group members contribute what they know.
- *The whole goes beyond the sum of the parts: synergy.* The give and take of group discussion might sift information and perspectives in a way that leads the group to a good solution to a problem, one in which the whole is actually more than the sum of its parts. In such cases, deliberation is, at the very least, an ambitious form of information aggregation, one in which the exchange of views leads to a creative answer or solution. If a group is seeking to improve the design of an automobile, a tennis racquet, or cell phone – or the right response to a threat to national security -- the exchange of ideas might well produce creative solutions that go far beyond any simple aggregation of what group members thought before they started to talk.

To what extent do these mechanisms work in practice? Two points are entirely clear. First, group members *tend to become far more confident of their judgments after*

*they speak with one another.*² A significant effect of group interactions is a greater sense that one's postdeliberation conclusion is correct—whether it actually is or not. Corroboration by others increases confidence in one's judgments.³ Second, deliberation *usually reduces variance*. After talking together, group members tend to come into accord with one another. It follows that members of deliberating groups will usually converge on a position on which members have a great deal of confidence. This is not disturbing if that position is also likely to be correct—but if it is not, then many group members will end up sharing a view in which they firmly believe, but which turns out to be wrong (a most unfortunate and sometimes quite destructive situation).

Unfortunately, there is no systematic evidence that deliberating groups will usually succeed in aggregating the information held by their members. With respect to questions with definite answers, deliberating groups tend to do about as well as or slightly better than their average members, but not as well as their best members.⁴ Hence, it is false to say that group members usually end up deferring to their internal specialists. Truth does not win out; the most that can be said is that under some conditions, the group will converge on the truth if the truth begins with “at least some initial support” within the group when the task has “a demonstrably correct answer.”⁵ (Recall the potential for group success with a eureka problem.) Note here that when a group outperforms most of its individual members, it is generally because the issue is one on which a particular answer can be shown, to the satisfaction of all or most, to be right; and that even in that condition, the group might not do well if the demonstrably correct solution lacks significant support at the outset.

² Chip Heath and Rich Gonzalez, “Interaction with Others Increases Decision Confidence but Not Decision Quality: Evidence against Information Collection Views of Interactive Decision Making,” *Organizational Behavior and Human Decision Processes* 61 (1995): 305.

³ See Robert Baron et al., “Social Corroboration and Opinion Extremity,” *Journal of Experimental Social Psychology* 32 (1996): 537.

⁴ Daniel Gigone and Reid Hastie, “Proper Analysis of the Accuracy of Group Judgments,” *Psychological Bulletin* 121 (1997): 149, 161; Reid Hastie, “Experimental Evidence of Group Accuracy,” in *Information Pooling and Group Decision Making*, ed. Bernard Grofman, Guillermo Owen, et al. (Greenwich, Conn.: JAI Press, 1986), 129.

⁵ Robert J. MacCoun, “Comparing Micro and Macro Rationality,” in *Judgments, Decisions, and Public Policy*, ed. Rajeev Gowda and Jeffrey Fox (Cambridge: Cambridge University Press, 2002), 116, 121.

In general, simple majority schemes do fairly well at predicting group judgments for many decision tasks. More generally, depending on the task, the appropriate statistical measure of the central predeliberation tendency – mean, median, or mode – is a good description of the group’s integration process. It follows that if the majority is wrong, the group will be wrong as well.⁶ With experts, the same general conclusion holds. Thus a “structured approach for combining independent forecasts is invariably more accurate” than “traditional group meetings,” which do “not use information efficiently.”⁷

Two Sources of Self-Silencing

There are two reasons why exposure to the views of others might lead people not to share what they know. The first involves the information conveyed by what other people say and do. If most group members believe that X is true, there is reason to believe that X is in fact true, and that reason might seem to outweigh the purely private reason a particular group member has to believe that X is false. If other group members share a particular belief, isolated or minority members might not speak out, deferring to the informational signal given by the statements of others. In a business or law firm, for example, members’ optimism about the likely success of some plan or option might lead skeptics to silence themselves on the ground that their own judgments must be ill-informed or wrong.

Not surprisingly, the strength of the informational signal will depend on the number and nature of the people who are giving it. If the group contains one or more people who are well-known to be authorities or who otherwise command respect, then other group members are likely to silence themselves out of deference to them. People are particularly averse to being sole dissenters.⁸ If all but one person in a deliberating group has said that X is true, then the remaining member is likely to agree X is true, even to the point of ignoring the evidence of his own senses. An evident danger here is that groups will be deprived of important information.

⁶ Ibid.

⁷ J. Scott Armstrong, “Combining Forecasts,” in *Principles of Forecasting*, ed. J. Scott Armstrong (Boston: Kluwer Academic, 2001), 433.

⁸ See Cass R. Sunstein, *Why Societies Need Dissent* (Cambridge: Harvard University Press, 2003).

The second reason for self-silencing involves social influences. 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. Even in societies and organizations that are strongly committed to freedom, those who defy the dominant position within the group might incur a form of disapproval that will lead them to be less trusted, liked, and respected in the future. Here, too, people are inevitably affected by the number and nature of those with the majority position. A large majority will impose more social pressure than a small one. And 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.

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

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

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

In addition, people show a strong tendency to self-edit when contributing information to a group discussion. When people have information that is in tension with their own conclusion, they tend not to share that information with others. One study of over 500 mock jury deliberations virtually never observed a juror contributing information to discussion that contradicted his or her currently preferred verdict.⁹ Sometimes this form of self-editing occurs as part of a deliberate strategy to persuade: It would be counter to my own interests to provide information that contradicts my own conclusions. But this same result can occur when the speaker has the best of intentions. Group members are simply not motivated to share information that they have already discounted in their personal judgment process – not for instrumental or strategic reasons, but because they truly believe that it would be counterproductive to share information that they deem false, misleading, or invalid.

Both informational pressure and social influences help to explain the finding that in a deliberating group, those in a minority position often silence themselves or otherwise have disproportionately little weight. There is a more particular finding: Members of low-status groups—less-educated people, African-Americans, sometimes women—speak less and carry less influence within deliberating groups than their higher-status peers.¹⁰ Both informational influence and social pressures, likely to be especially strong for low-status members, contribute to this result. The unfortunate consequence can be a loss of information to the group as a whole, in a way that ensures that deliberating groups do far less well than they would if only they could aggregate the information held by group members. Consider the finding that one of the most common defensive tactics for small coalition or solo hold-out jurors is to refuse to discuss the case further with other jurors.¹¹

More generally, a comprehensive study has demonstrated that majority pressures can be powerful even for factual questions to which some people know the right

⁹ See Reid Hastie et al., *Inside the Jury* (Cambridge: Harvard University Press, 1983).

¹⁰ Caryn Christenson and Ann Abbott, “Team Medical Decision Making,” in *Decision Making in Health Care*, ed. Gretchen Chapman and Frank Sonnenberg (New York: Cambridge University Press, 2000), 267, 273–76.

¹¹ Hastie et al., *supra* note.

answers.¹² The study involved twelve hundred people, forming groups of six, five, and four members. Individuals were asked true-false questions involving art, poetry, public opinion, geography, economics, and politics. They were then asked to assemble into groups, which discussed the questions and produced answers. The majority played a substantial role in determining each group's answers. The truth played a role, too, but a lesser one. If a majority of individuals in the group gave the right answer, the group's decision moved toward the majority in 79 percent of the cases. If a majority of individuals in the group gave the wrong answer, the group's decision nonetheless moved toward the majority in 56 percent of the cases.

Hence, the truth did have an influence—79 percent is higher than 56 percent—but the majority's judgment was the dominant one. And because the majority was influential even when wrong, the average group decision was right only slightly more often than the average individual decision (66 percent versus 62 percent). What is most important is that groups did not perform as well as they would have if they had properly aggregated the information that group members had.

Let us now turn to the four sets of deliberative failures.

Deliberative Failure 1: Amplification of Cognitive Errors

It is well known that individuals do not always process information well. They use heuristics that lead them to predictable errors; they are also subject to identifiable biases, which also produce errors.¹³

For example, people err because they use the *availability heuristic* to answer difficult questions about probability. How likely is a terrorist attack, a hurricane, a traffic jam, an accident from a nuclear power plant, a case of venereal disease? When people use the availability heuristic, they answer a question of probability by asking whether examples come readily to mind.¹⁴ The point very much bears on private and public responses to

¹²Robert L. Thorndike, "The Effect of Discussion upon the Correctness of Group Decisions: When the Factor of Majority Influence Is Allowed For," *Journal of Social Psychology* 9 (1938): 343.

¹³For an overview, see Thomas Gilovich, Dale Griffin, and Daniel Kahneman, *Heuristics and Biases: The Psychology of Intuitive Judgment* (New York: Cambridge University Press, 2002).

¹⁴See Amos Tversky & Daniel Kahneman, *Availability: A Heuristic For Judging Frequency*

risks -- suggesting, for example, that people will be especially responsive to the dangers of AIDS, crime, earthquakes, and nuclear power plant accidents if examples are easy to recall. The point also explains some of the sources of discrimination on the basis of race, sex, age, and disability. If it is easy to bring to mind cases in which a female employee quit work to care for her family, sex discrimination is more likely; so too if it is easy to think of cases in which African-American employees performed poorly.

In this way, *familiarity* can affect the availability of instances. But *salience* is important as well. A terrorist attack on television will be highly salient to viewers, and will have a greater impact than a report about the attack in the newspaper.¹⁵ Similarly, earlier events will have a smaller impact than later ones. The point helps explain much behavior. For example, whether people will buy insurance for natural disasters is greatly affected by recent experiences.¹⁶ In the aftermath of an earthquake, people become far readier to buy insurance for earthquakes – but their readiness to do so declines steadily from that point, as vivid memories recede. Use of the availability heuristic is not irrational, but it can easily lead to serious errors of fact. After the 2005 disaster produced by hurricane Katrina in the United States, it was predictable that significant steps would be taken to prepare for hurricanes – and also predictable that before that disaster, such steps would be quite inadequate.

Most people also 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).¹⁷ The representativeness heuristic helps explain what Paul Rozin and Carol Nemeroff have called “sympathetic magical thinking,” including the beliefs that some objects have contagious properties, and that causes resemble their effects.¹⁸ The representativeness heuristic often works well, but it can also lead to severe blunders.

and Probability, 5 *Cognitive Psychol.* 207, 208 (1973) (discussing availability heuristic).

¹⁵ Paul Slovic, *The Perception of Risk* 37-48 (London: Earthscan Publications, 2000).

¹⁶ *Id.* at 40

¹⁷ Amos Tversky and Daniel Kahneman, “Judgment under Uncertainty: Heuristics and Biases,” in *Judgment under Uncertainty: Heuristics and Biases*, ed. Daniel Kahneman, Paul Slovic, and Amos Tversky (Cambridge: Cambridge University Press), 1982, 3.

¹⁸ Paul Rozin and Carol Nemeroff, “Sympathetic Magical Thinking: The Contagion and Similarity ‘Heuristics,’” in Gilovich, Griffin, and Kahneman, *Heuristics and Biases* p. 201

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

Consider some key findings. If individual jurors are biased because of pretrial publicity that misleadingly implicates the defendant, or even because of the defendant’s unappealing physical appearance, juries are likely to amplify rather than correct those biases.¹⁹ Groups have been found to amplify, rather than to attenuate, reliance on the representativeness heuristic;²⁰ to reflect even larger framing effects than individuals;²¹ to show more overconfidence than group members;²² to be more affected by the biasing effect of spurious arguments from lawyers;²³ to be more susceptible to the “sunk cost fallacy”;²⁴ and to be more subject to choice-rank preference reversals.²⁵ 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.²⁶ In addition, groups demonstrate essentially the same level of reliance on the availability heuristic, even when use of that heuristic leads to clear errors.²⁷

¹⁹ MacCoun, “Comparing Micro and Macro Rationality,” 116, 121.

²⁰ Mark F. Stasson et al., “Group Consensus Approaches in Cognitive Bias Tasks,” *Japanese Psychological Research* 30 (1988): 68.

²¹ See Norbert L. Kerr et al., “Bias in Judgment: Comparing Individuals and Groups,” *Psychology Review* 103 (1996): 687, 689, 691–93.

²² Janet A. Sniezek and Rebecca A. Henry, “Accuracy and Confidence in Group Judgment,” *Organizational Behavior and Human Decision Processes* 43 (1989): 1. This finding very much bears on excessive risk-taking, including in the context of making war. See Dominic Johnson, *Overconfidence and War*, 180–83.

²³ Edward L. Schumann and W. C. Thompson, “Effects of Attorney’s Arguments on Jurors’ Use of Statistical Evidence” (unpublished manuscript, 1989).

²⁴ Glen Whyte, “Escalating Commitment in Individual and Group Decision Making,” *Organizational Behavior and Human Decision Processes* 54 (1993): 430.

²⁵ James W. Gentry and John C. Mowen, “Investigation of the Preference Reversal Phenomenon in a New Product Introduction Task,” *Journal of Applied Psychology* 65 (1980): 715; Julie R. Irwin and James H. Davis, “Choice/Matching Preference Reversals in Groups,” *Organizational Behavior and Human Decision Processes* 64 (1995): 325.

²⁶ Whyte, “Escalating Commitment,” 430.

²⁷ Stasson et al., “Group Consensus Processes,” 68.

Here is an especially disturbing finding, one with great relevance to group behavior in both business and politics: *Groups are more likely than individuals to escalate their commitment to a course of action that is failing -- and all the more so if members identify strongly with the groups of which they are a part.*²⁸ There is a clue here about why companies, states, and even nations often continue with projects and plans that are clearly going awry. If a company is marketing a product that is selling poorly, it may well continue on its misguided course simply because of group dynamics. (Enron is a likely example.) So too with a nation whose economic policy, or approach to foreign affairs, is hurting its citizens.

Why are individual cognitive errors so often amplified at the group level? Informational pressures and social influences are unquestionably at work. Suppose that most members of a group are prone to make certain errors. If the majority makes those errors, then most people will see others making the same errors. What they see will convey information about what is right. Those who are not specialists are likely to think: If most people make the same errors, maybe they are not errors at all. Social influences also play a role. If most group members make errors, others also might make them simply in order not to seem disagreeable or foolish. When groups do not correct the blunders of their members, this is a large reason.

To be sure, there is some evidence that deliberating groups can attenuate certain biases. We have seen that for eureka problems, groups often do well, even if individual members begin with an answer produced by some kind of bias. Groups are also likely to do better than the average individual when members are subject to “egocentric bias” – the bias that leads most of us to think that other people think and act as we do. Most people believe that their tastes and preferences are typical. When asked what percentage of other people watch television on Saturday night, enjoy Frank Sinatra, favor a particular political party, or believe that the latest Brad Pitt movie will win the Oscar, most of us show a bias in the direction that we ourselves favor. But in groups with diverse views, people quickly learn that their own position is not universally held, and hence the bias is reduced. In these cases, group deliberation supplies an important corrective.

Or consider the hindsight bias — people’s tendency to believe, falsely but with the

²⁸ Stasser & Dietz-Uhler, *supra* note, at 48.

benefit of hindsight, that they would have accurately predicted the outcome of an event (an accident, a natural disaster, an illness, a change in the stock market). Compared to individuals, groups are slightly less susceptible to hindsight bias.²⁹ Apparently group members who are not susceptible to that bias are able to persuade others that it is indeed a bias.

But the broader point is that with group discussion, individual errors are usually propagated rather than eliminated, and amplification of those errors is quite likely. When individuals show a high degree of bias, groups are likely to be more biased, not less biased, than their median or average member.

Deliberative Failure 2: Cascades

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

Informational Cascades. Cascades need not involve deliberation, but deliberative processes often involve cascades. 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, and it often decides badly.³⁰

To see how informational cascades work, imagine a deliberating group that is deciding whether to authorize some new venture.³¹ Let us also assume that group members are announcing their views in sequence, in a temporal queue, and that each member knows his place in that queue. Every member has some private information

²⁹ See generally Dagmar Stahlberg et al., *We Knew It All Along: Hindsight Bias in Groups*, 63 *Organizational Behav. & Hum. Decision Processes* 46 (1995).

³⁰ The literature here is vast. For a superb summary, see David Hirschleifer, "The Blind Leading the Blind," in *The New Economics of Human Behavior*, ed. Marianno Tommasi and Kathryn Ierulli 188 (Cambridge: Cambridge University Press, 1995)

³¹ We draw here on David Hirschleifer, "The Blind Leading the Blind," in *The New Economics of Human Behavior*, ed. Marianno Tommasi and Kathryn Ierulli (Cambridge: Cambridge University Press, 1995), 188, 193–94.

about what should be done. But each also attends, reasonably enough, to the judgments of others. Mr. Andrews is the first to speak. He suggests that the venture should be authorized. Ms. Barnes now knows Andrews's judgment; it is clear that she, too, should vote in favor of the venture if she agrees independently with Andrews. But if her independent judgment is otherwise, she would—if she trusts Andrews no more and no less than she trusts herself—be indifferent about what to do and might simply flip a coin.

Now turn to a third member, Mr. Carlton. Suppose that both Andrews and Barnes have argued in favor of the venture but that Carlton's own information, though inconclusive, suggests that the venture is a terrible idea. In that event, Carlton might well ignore what he knows and follow Andrews and Barnes. It is likely in these circumstances that both Andrews and Barnes had reasons for their conclusion, and unless Carlton thinks that his own information is better than theirs, he should follow their lead. If he does, Carlton is in a cascade. Now suppose that Carlton is acting in response to what Andrews and Barnes did, not on the basis of his own information, and that subsequent members know what Andrews, Barnes, and Carlton did. On reasonable assumptions, they will do exactly what Carlton did: favor the venture regardless of their private information (which, we are supposing, is relevant but inconclusive). This will happen even if Andrews initially blundered.³²

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

Note that there is nothing irrational, from the point of view of each group member, in following the statements or actions of predecessors; members are reacting rationally to the informational signals they receive. Note as well that even when people

³² *Ibid.*, 195.

are motivated to discuss an issue vigorously, they often focus on developing the statement of their opinions, and not on the evidence that led to those opinions in the first place.

Cascades often occur in the real world within deliberating groups;³³ they are easy to create in the laboratory. The simplest experiment asked subjects to guess whether the experiment was using urn A, which contained two red balls and one white, or urn B, which contained two white balls and one red.³⁴ Subjects could earn \$2.00 for a correct decision, and hence an economic incentive favored correct individual decisions. In each period, the contents of the chosen urn were emptied into a container. A randomly selected subject was asked to make one (and only one) private draw of a ball in each round. The subject recorded the color of that draw on an answer sheet and his own decision about which urn was involved. The subject did not announce his draw to the group, but he did announce his own decision to everyone. Then the urn was passed to the next subject for his own private draw, which again was not disclosed, and his own decision about the urn, which again was disclosed. This process continued until all subjects had made draws and decisions. At that time, the experimenter announced the actual urn used. If the subject had picked the urn only on the basis of his private information, he would have been right 66.7 percent of the time. The point of the experiment was to see whether people will decide to ignore their own draw in the face of conflicting announcements by predecessors—and to explore whether such decisions will lead to cascades and errors.

In the experiment, cascades often developed and often produced errors. After a number of individual judgments were revealed, people sometimes announced decisions that were inconsistent with their private draws, but that fit with the majority of previous announcements.³⁵ More than 77 percent of “rounds” resulted in cascades, and 15 percent of private announcements did not reveal a “private signal,” that is, the information provided by people’s own draws. Consider cases in which one person’s draw (say, red) contradicted the announcement of his predecessor (say, urn B). In such cases, the second

³³ See *ibid*; also see Sunstein, *Why Societies Need Dissent*.

³⁴ See Lisa Anderson and Charles Holt, “Information Cascades in the Laboratory,” *American Economic Review* 87 (1997): 847.

³⁵ See Angela Hung and Charles Plott, “Information Cascades: Replication and an Extension to Majority Rule and Conformity-Rewarding Institutions,” *American Economic Review* 91 (2001): 1508, 1515.

announcement nonetheless matched the first about 11 percent 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. Of note, the majority of decisions were rationally based on the available information³⁶—but erroneous cascades nonetheless developed. Table 4-1 shows an example of a cascade that produced an inaccurate outcome (the urn used was B).³⁷

Table 4-1. An Informational Cascade

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

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

What is noteworthy here, of course, is that the total amount of private information—four whites and two reds—justified the correct judgment, which was in favor of urn B. But the existence of two early signals, producing rational but incorrect judgments, led everyone else to fall in line. “Initial misrepresentative signals start a chain of incorrect decisions that is not broken by more representative signals received later.”³⁸ This result maps directly onto real-world decisions by deliberating groups, in which people fail to disclose what they know, to the detriment of the group as a whole.

Reputational Cascades. In a reputational cascade, people think they know what is right, or what is likely to be right, but they nonetheless go along with the crowd in order to maintain the good opinion of others. The problem is not that group members are influenced by the information contained in the statements of their predecessors, but they people do not want to face public disapproval. “Political correctness” can be found in many places; it is hardly limited to left-leaning institutions of education.

³⁶ Thus, 72 percent of subjects followed Bayes’s rule in Lisa Anderson and Charles Holt, “Information Cascades in the Laboratory,” *American Economic Review* 87 (1997): 847, and 64 percent in Marc Willinger and Anthony Ziegelmeyer, “Are More Informed Agents Able to Shatter Information Cascades in the Lab?” in *The Economics of Networks: Interaction and Behaviours*, ed. Patrick Cohendet et al. (New York: Springer, 1998), 291, 304.

³⁷ See Willinger and Ziegelmeyer, “Are More Informed Agents,” 291.

³⁸ See Anderson and Holt, “Information Cascades in the Laboratory,” 847.

Suppose Albert suggests that global warming is a serious problem and that Barbara concurs with Albert, not because she actually thinks that Albert is right, but because she does not wish to seem, to Albert, ignorant or indifferent to environmental protection. If Albert and Barbara seem to agree that global warming is a serious problem, Cynthia not only might not contradict them publicly, but also might even appear to share their judgment, not because she believes that judgment to be correct, but because she does not want to face their hostility or lose their good opinion.

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

The possibility of reputational cascades is demonstrated by an ingenious variation on the urn experiment outlined above.³⁹ In this experiment, people were paid \$0.25 for a correct decision, but \$0.75 for a decision that matched the decision of the majority of the group. There were punishments for incorrect and nonconforming answers as well. If people made an incorrect decision, they lost \$0.25; if their decision failed to match the group's decision, they lost \$0.75.

In this experiment, cascades appeared almost all of the time! No fewer than 96.7 percent of rounds resulted in cascades, and 35.3 percent of people's announcements did not match their private signal, that is, the signal given by their own draw. And when the draw of a subsequent person contradicted the announcement of the predecessor, 72.2 percent of people matched the first announcement. Consider, as a dramatic illustration, table 4-2, which shows this period of the experiment (the actual urn was B):⁴⁰

³⁹ See Hung and Plott, "Information Cascades," 1515–17.

⁴⁰ *Ibid.*, 1516.

Table 4-2. Conformity and Cascades

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

Source: Hung and Plott, "Information Cascades."

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

Availability cascades. Thus far we have assumed that people are, in a sense, completely rational. They are listening to one another and paying attention to the informational signals given by their predecessors. True, they care about their reputations, but there is nothing irrational about that.

As we have noted, however, people use heuristics, which can lead them astray, and they are also subject to biases. For purposes of understanding how cascade effects can go wrong, the most important heuristic involves availability. When a particular event is salient, it can lead to *availability cascades*, which arise as an event becomes cognitively salient, and spreads rapidly from one person to another, eventually producing a widespread belief.⁴¹

In the area of risk, availability cascades are common. A particular event – involving shark attacks, a dangerous pesticide, an abandoned hazardous waste dump, a nuclear power accident, an act of terrorism – may become well-known in the public mind, even iconic, and if so, it will alter public perceptions of a process, product, or activity. In business, availability cascades are also common. A particular event, a particular success, or a particular failure may spread like wildfire within or across firms, leading to a judgment about other apparently similar events or products. If a movie (Star Wars?), a television show (Survivor?), or a book (Freakonomics?), does well, businesses will react strongly, looking for a proposal that seems relevantly similar.

Of course the underlying judgments might be correct. But recall that the availability heuristic is unreliable: People might know of an instance in which a risk came

⁴¹ See Timur Kuran and Cass R. Sunstein, Availability Cascades and Risk Regulation, 51 Stan L Rev 683 (1988).

to fruition, but the instance may not be representative. Pesticides might be safe even if a particular pesticide is not; a well-publicized event involving an abandoned hazardous waste dump, such as Love Canal, may suggest that abandoned hazardous waste dumps are far more dangerous than they are in fact; nuclear power might be safer than the alternatives, even if people's thoughts are dominated by images of the Chernobyl disaster or the accident at Three Mile Island.

Similar errors infect firms. A company might believe that a particular product is likely to succeed, because it is like some other product that has succeeded, even though the chance of success is low. Television broadcasters often blunder for this reason, thinking that because some science fiction show has succeeded, another one will too. Or a company might draw an undue inference from some failure, thinking that similar plans will fail too. Availability cascades make businesses and government unrealistically optimistic about some possibilities and unrealistically pessimistic about others.

Deliberative Failure 3: Group Polarization

There are clear links between social cascades and the well-established phenomenon of group polarization, by which *members of a deliberating group end up adopting a more extreme version of the position toward which they tended before deliberation began.*⁴² The problem is especially severe for groups of like-minded people, who typically end up in more extreme positions as a result of deliberation.

Group polarization is the typical pattern with deliberating groups, and it has been found in hundreds of studies involving more than a dozen countries, including the United States, France, Afghanistan, and Germany.⁴³ For example, those who disapprove of the United States and are suspicious of its intentions will increase their disapproval and suspicion if they exchange points of view. Indeed, there is specific evidence of the latter phenomenon among citizens of France.⁴⁴

The original psychological experiments on the effects of deliberation are relevant to businesses and governments alike. They involved risk-taking behavior, with a

⁴² Roger Brown, *Social Psychology: The Second Edition* (New York, N.Y.: Free Press, 1986), 206–7.

⁴³ *Ibid.*, 204.

⁴⁴ *Ibid.*, 224.

demonstration that risk-inclined people become still more risk-inclined after they deliberate with one another.⁴⁵ Risky decisions include taking a new job, investing in a foreign country, escaping from a prisoner-of-war camp, or running for political office. With respect to many such decisions, members of deliberating groups became significantly more risk-inclined after a brief period of collective discussion. On the basis of this evidence, it became standard to believe that deliberation produced a systematic “risky shift.”

Later studies drew this conclusion into serious question. On many of the same questions on which Americans displayed a risky shift, Taiwanese subjects showed a “cautious shift.” Deliberation led citizens of Taiwan to become significantly less risk-inclined than they were before they started to talk. Among American subjects, deliberation sometimes produced a cautious shift as well, as risk-averse people became more averse to certain risks after they talked with one another. The principal examples of cautious shifts involved the decision whether to marry and the decision whether to board a plane despite severe abdominal pain, possibly requiring medical attention. In these cases, the members of deliberating groups shifted not toward risk but toward greater caution.⁴⁶

A straightforward interpretation turned out to be able to reconcile these competing findings: the pre-deliberation median is the best predictor of the direction of the shift.⁴⁷ When group members are disposed toward risk-taking, a risky shift is observed. Where members are disposed toward caution, a cautious shift is observed. Thus, for example, the striking difference between American and Taiwanese subjects is a product of a difference in the pre-deliberation medians of the different groups on the relevant questions. Thus the risky shift and the cautious shift are both subsumed under the rubric of group polarization.

Group polarization occurs for matters of fact as well as issues of value, though it is easier to demonstrate the latter. If the question is whether a terrorist attack will occur in the United States in the next year, group polarization will not be easy to test, simply

⁴⁵ For a clear discussion, see *id.*

⁴⁶ S. Moscovici and M. Zavalloni, The group as a polarizer of attitudes 12. *Journal of Personality and Social Psychology* 125 (1969).

⁴⁷ *Id.*; Brown, *supra* note.

because the answer is either yes or no, and it is not simple to demonstrate greater extremism in binary choices. But suppose that people are asked, on a bounded scale of zero to eight, how likely it is that a terrorist attack will occur in the United States in the next year, with zero indicating “zero probability,” eight indicating “absolutely certain,” seven indicating “overwhelmingly likely,” six “more probable than not,” and five “fifty-fifty.” In that event, the answers from a deliberating group will tend to reveal group polarization, as people move toward more extreme points on the scale depending on their initial median point. If the predeliberation median is five, the group judgment will usually be six; if the predeliberation median is three, the group judgment will usually be two.⁴⁸ Note here that even 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.⁴⁹ Juries polarize as well.⁵⁰

Why does group polarization occur? There are three reasons.⁵¹ The first and most important involves the now-familiar idea of informational influence, but in a distinctive form. People respond to the arguments made by other people—and the “argument pool” in any group with some predisposition in one direction will inevitably be skewed toward that predisposition. As a statistical matter, the arguments favoring the initial position will be more numerous than those pointing in the other direction. Individuals will have heard of some, but not all, of the arguments that emerge from group deliberation. As a result of the relevant arguments, deliberation will lead people toward a more extreme point in line with what group members initially believed.

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

⁴⁸ Ibid.

⁴⁹ Cass R. Sunstein et al., *Are Judges Political? An Empirical Investigation of the Federal Judiciary* (Washington, DC: Brookings, 2005).

⁵⁰ See David Schkade et al., “Deliberating about Dollars: The Severity Shift,” *Columbia Law Review* 100 (2000): 101.

⁵¹ Brown, *Social Psychology*, 200–45. An effort to systematize some of these points can be found in Edward Glaeser and Cass R. Sunstein, *Extremism and Social Learning*, 1 *J Legal Analysis* (forthcoming 2008).

of the dominant position in order to hold onto their preserved self-presentation. They shift accordingly.⁵²

The third explanation of group polarization stresses the close links among confidence, extremism, and corroboration by others.⁵³ As people gain confidence, they usually become more extreme in their beliefs. Agreement from others tends to increase confidence, and for this reason like-minded people, having deliberated with one another, become more sure that they are right and thus more extreme. In many contexts, people's opinions become more extreme simply because their views have been corroborated, and because they become more confident after learning that others share their views.⁵⁴ Many minds can badly blunder in this way.

A great deal of work suggests that group polarization is heightened when people have a sense of shared identity, and this point helps to suggest yet another explanation of polarization.⁵⁵ People may polarize because they are attempting to conform to the position that they see as typical within their own group. If their group's identity is made especially salient, the in-group norms are likely to become more extreme.⁵⁶ There is a further point. If arguments come from a member of an in-group, they are especially likely to be persuasive. Such arguments are more likely to seem right. It is also reasonable to think that people would fear the social pressures that come from rejecting what an in-group member has to say. By contrast, the views of out-group members have far less force and might even be irrelevant.⁵⁷ Should you really listen to someone who belongs to a group that you believe to be systematically wrong? The reputational pressure will also be reduced: Do you really have to worry about rejecting the views of someone usually thought to be wrong by people like you? If a self-identified conservative hears arguments

⁵² Brown, *supra* note. It has similarly been suggested that majorities are especially potent because people do not want to incur the wrath, or lose the favor, of large numbers of others, and that when minorities have influence, it is because they produce genuine attitudinal change. See Baron et al., "Social Corroboration," 82.

⁵³ See Baron et al., *supra* note, at 557–59 (showing that corroboration increases confidence and hence extremism).

⁵⁴ *Id.* at 541, 546–47, 557 (concluding that corroboration of one's views has effects on opinion extremity).

⁵⁵ See Brown, *supra* note, at 209–11; Turner, *supra* note, at 159–70; Cooper et al., *supra* note, at 259, 269–70.

⁵⁶ Brown, *supra* note, at 210.

⁵⁷ *Id.* at 211; Cooper et al., *supra* note, at 269.

from someone who is known to be left of center, those arguments are likely to fall on deaf ears. The clear lesson is that when a group is highly cohesive, and when its members are closely identified with it, polarization is especially likely -- and likely to be especially large.

Does group polarization lead 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 tendencies. If the group is leaning toward the right answer, polarization might lead them directly to the truth. It follows that group polarization might not produce errors or failures at all. But there are no guarantees here. As a result of the relevant influences, some people will fail to disclose what they know. 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 logical errors, group processes lead to more errors rather than fewer. This is polarization in action, and it produces large blunders.

Deliberative Failure, 4: Over-weighting Common Knowledge

The last failure is in many respects the most interesting. Suppose that group members have a great deal of information—enough to produce the unambiguously right outcome if that information is properly aggregated. Even if this is so, an obvious problem is that groups will not perform well if they emphasize shared information and slight information that is held by one or a few members. Unfortunately, countless studies demonstrate that this regrettable result is highly likely.⁵⁸

“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.⁵⁹ The most obvious explanation

⁵⁸ Garold Stasser and William Titus, “Hidden Profiles: A Brief History,” *Psychological Inquiry* 14 (2003): 304.

⁵⁹ Daniel Gigone and Reid Hastie, “The Common Knowledge Effect: Information Sharing and Group Judgments,” *Journal of Personality and Social Psychology* 65 (1993): 959.

of the effect is the simple fact that as a statistical matter, common knowledge is more likely to be communicated to the group; but social influences play a role as well.

Hidden Profiles. Consider a study of serious errors within working groups, both face-to-face and online.⁶⁰ The purpose of the study was to see how groups might collaborate to make personnel decisions. Resumes for three candidates applying for a marketing manager position were placed before group members. The attributes of the candidates were rigged by the experimenters so that one applicant was clearly the best for the job described. Packets of information were given to subjects, each containing a subset of information from the resumes, so that each group member had only part of the relevant information. The groups consisted of three people, some operating face-to-face, some operating online.

Almost none of the deliberating groups made what was conspicuously the right choice! The reason is simple: They failed to share information in a way that would permit the group to make that choice. Members tended to share positive information about the winning candidate and negative information about the losers. They suppressed negative information about the winner and positive information about the losers. Hence, their statements served to “reinforce the march toward group consensus rather than add complications and fuel debate.”⁶¹

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.⁶² In the first condition, each member of the four-person groups was given most of the relevant information (66 percent of the information about each candidate). In that condition, 67 percent of group members favored candidate A before discussion and 85 percent after discussion.⁶³ This is a clear example of appropriate aggregation of information. Groups significantly outperformed individuals, apparently because of the

⁶⁰ See Ross Hightower and Lutfus Sayeed, “The Impact of Computer-Mediated Communication Systems on Biased Group Discussion,” *Computers in Human Behavior* 11 (1995): 33.

⁶¹ Patricia Wallace, *The Psychology of the Internet* (Cambridge: Cambridge University Press, 1999), 82.

⁶² See Garold Stasser and William Titus, “Pooling of Unshared Information in Group Decision Making: Biased Information Sampling during Discussion,” *Journal of Personality and Social Psychology* 48 (1985): 1467.

⁶³ *Ibid.*, 1473; see also Stasser and Titus, “Hidden Profiles,” 304.

exchange of information and reasons. Here, then, is a clear illustration of the possibility that groups can aggregate what members know in a way that produces sensible outcomes.

In the second condition, by contrast, the information that favored candidate A was parceled out to various members of the group so that only 33 percent of information about each candidate was shared. As the condition was designed, the shared information favored two unambiguously inferior candidates, B and C; but if the unshared information emerged through discussion, and were taken seriously, candidate A would be chosen. In that condition, less than 25 percent of group members favored candidate A before discussion, a natural product of the initial distribution of information. But (and this is the key result) that number actually *fell* after discussion, simply because the shared information had disproportionate influence on group members.⁶⁴ In other words, groups did worse, not better, than individuals when the key information was distributed selectively. In those conditions, the commonly held information was far more influential than the distributed information, to the detriment of the group's ultimate decision.

From this and many similar studies, the general conclusion is that when "the balance of unshared information opposes the initial most popular position . . . the unshared information will tend to be omitted from discussion and, therefore, will have little effect on members' preferences during group discussion."⁶⁵ It follows that "group decisions and postgroup preferences reflect the initial preferences of group members even when the exchange of unshared information should have resulted in substantial shifts in opinion."⁶⁶ 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.⁶⁷ 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."⁶⁸

The Common-Knowledge Effect. These results are best understood as a consequence of the common-knowledge effect, by which information held by all group

⁶⁴ Stasser and Titus, "Pooling of Unshared Information," 1473.

⁶⁵ *Ibid.*, 1476.

⁶⁶ *Ibid.*

⁶⁷ *Ibid.*

⁶⁸ Stasser and Titus, "Hidden Profiles," 305.

members has far more influence on group judgments than information held by one member or a few.⁶⁹ 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.”⁷⁰ 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.”⁷¹

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 groups.⁷² The more shared information is (that 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.

Most of the hidden profile experiments involve volunteer participants from college courses. It is natural to ask: Would the same results be found in the real world? Affirmative evidence comes from a hiring exercise involving high-level executives.⁷³ In this study, no experimenter controlled information about the various candidates. Instead the executives’ knowledge arose naturally from their own information searches. As a result of those searches, some information was known to all; some was partially shared; and some was uniquely held. Even with high-level executives, common information had a disproportionately large impact on discussions and conclusions. Disproportionately little weight was given to valuable information held by one person or a few.

The same study offers an additional finding of considerable importance. Some group members are “cognitively central,”⁷⁴ in the sense that their knowledge is shared

⁶⁹ See Daniel Gigone and Reid Hastie, “The Common Knowledge Effect: Information Sharing and Group Judgments,” *Journal of Personality and Social Psychology* 65 (1993): 959.

⁷⁰ *Ibid.*, 960.

⁷¹ *Ibid.*, 973.

⁷² *Id.*

⁷³ Susanne Abele et al., *Information and Cognitive Centrality* (May 2005), available at www.erim.eur.nl

⁷⁴ See T. Kameda et al., *Centrality in Sociocognitive Networks and Social Influence: An Illustration in a Group Decision-Making Context*, 73 *J Personality and Social Psych* 296 (1997).

with many other group members. What they know, other people know as well. A cognitively central group member is defined as one who possesses information in common with all or most group members. But other group members are “cognitively peripheral,” in the sense that their own information is uniquely held; what they know is known by no one else. Well-functioning groups need to take advantage of cognitively peripheral people. But it turns out that cognitively central people usually have a disproportionate influence in discussion – and they also show higher levels of participation in group deliberations. By contrast, cognitively peripheral people end up having little influence and do not much participate.

A simple explanation for these results is that group members prefer to hear information that is commonly held – and prefer to hear people who have such information. Cognitively central people also have high levels of credibility; cognitively peripheral people have corresponding low levels. And indeed, the executives who were cognitively central ended up having a greater influence on the final report. The general conclusion is that when some group members count for more than others, it is often because they know what everyone else knows. Unfortunately, more peripheral people, with unique knowledge, are sometimes the ones group members most need to hear.

As might be expected, a group’s focus on shared information increases with the size of the group.⁷⁵ For this reason, groups can go very badly wrong even if they are fairly large. In a study designed to test judgments about candidates for office, involving both three-person and six-person groups, all group 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.”⁷⁶

It follows that for very large groups – twelve people, thirty people, two hundred people – the effect of shared information will be compounded. And despite the failures of their deliberations, group members tended to be significantly more confident about their judgments after discussion – an especially alarming finding, because confidence and error

⁷⁵ See Garold Stasser et al., *Information Sampling in Structured and Unstructured Discussions of Three- and Six-Person Groups*, 57 *J. Personality & Soc. Psychol.* 67, 72–73 (1989).

⁷⁶ *Id.* at 78.

are a bad combination.

Hiding profiles. Why do hidden profiles remain hidden? The principal explanations build on the informational and social influence accounts emphasized thus far. When information is held by all or most group members, it is especially 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. There are two different points here. First, 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. Second, shared information, simply because it is shared, is more likely to be explored during group discussion.

Suppose that a team of five people is advising the President whether to embark on military action to combat a perceived threat. If each of the five people has information indicating that the use of military force would be successful, that information is more likely to emerge in group discussion than are separate parcels of information, individually held by each adviser, suggesting that the use of force would run into trouble. If the team of advisers stresses the information that is held in advance by each, that information will have a disproportionate influence on its ultimate decision. This is a simple statistical point.

But hidden profiles remain significantly more hidden than would be predicted by statistics. 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.”⁷⁷ Those in the 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. It follows that in a business, leaders are unlikely to receive the views of those who do not have much experience or respect, even if they know something important.

This finding suggests that group members, and especially lower status ones, are nervous about emphasizing information that most group members lack. Indeed, lower status members “are likely to drop unique information like a hot potato”—partly because of the difficulty of establishing its credibility and relevance, and partly because they risk

⁷⁷ Id.

the group's disapproval if they press a line of argument that others reject. Return here to the finding that cognitively peripheral members have little influence on the group. Those who have uniquely held information end up participating less than those who have shared information, and what they have to say is not much valued. In many deliberating groups, people who emphasize uniquely held information take an obvious social risk; and they know it.

This point creates a big problem for such groups, which will fail to get information that they need, and which can therefore be lead to error. Consider the finding that group members typically *underestimate* the performance of low-status members and typically *overestimate* the performance of high-status members, in a way that gives high-status members a degree of deference that is not warranted by reality.⁷⁸

In the same vein, those who discuss shared information obtain rewards in the form of an enhanced sense of competence in the eyes of others—and in their own eyes as well.⁷⁹ Strange but true: If someone tells you something you already know, you are likely to like that person, and yourself, a bit better as a result. In face-to-face discussions and purely written tasks, people give higher ratings (in terms of knowledge, competence, and credibility) both to themselves and to others after receiving information that they knew already. It follows that someone with valuable, unshared information may do best to earn credibility “by telling others what they already know before telling them what they do not already know.”⁸⁰ The general problem is that deliberating groups often perform poorly because they fail to elicit information that could steer them in the right directions.

Conclusion

Our goal here has been to catalogue four common failures of deliberating groups. (1) Sometimes individual errors are amplified, not merely propagated, as a result of deliberation. (2) Groups fall victim to cascade effects, as the early speakers or actors ensure that people do not learn what is known by their successors. (3) Because of group

⁷⁸ Cecilia L. Ridgeway, *Social Status and Group Structure*, in *Group Processes*, *supra* note 93, at 352, 354 (collecting studies).

⁷⁹ See Gwen M. Wittenbaum et al., *Mutual Enhancement: Toward an Understanding of the Collective Preference for Shared Information*, 77 *J. Personality & Soc. Psychol.* 967, 967–78 (1999).

⁸⁰ Stasser & Titus, *Hidden Profiles*, *supra* note, at 311.

polarization, members of deliberating groups often end up in a more extreme position in line with their predeliberation tendencies. (4) In deliberation, shared information often dominates or crowds out unshared information, ensuring that groups do not learn what their members know.

As we have seen, informational signals and reputational pressure help to explain all four errors. A disturbing result is that groups often do not improve on, and sometimes do worse than, their average or median member. A crucial important task, for many deliberating groups, is to devise strategies that will overcome the relevant problems and increase the likelihood that group judgments will take advantage of the knowledge held by group members. We have not said anything here about how that task might be accomplished.⁸¹ Our goal has been instead to isolate the problems to which deliberating groups are subject, and to understand the mechanisms that threaten to make those problems so difficult to solve.

Readers with comments should address them to:

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

⁸¹ These issues will be explored in detail in Hastie and Sunstein, *Smart Groups*, supra note; some tentative thoughts can be found in Sunstein, *Infotopia*, supra note.

Chicago Working Papers in Law and Economics
(Second Series)

For a listing of papers 1–299 please go to Working Papers at
<http://www.law.uchicago.edu/Lawecon/index.html>

300. Adam B. Cox, The Temporal Dimension of Voting Rights (July 2006)
301. Adam B. Cox, Designing Redistricting Institutions (July 2006)
302. Cass R. Sunstein, Montreal vs. Kyoto: A Tale of Two Protocols (August 2006)
303. Kenneth W. Dam, Legal Institutions, Legal Origins, and Governance (August 2006)
304. Anup Malani and Eric A. Posner, The Case for For-Profit Charities (September 2006)
305. Douglas Lichtman, Irreparable Benefits (September 2006)
306. M. Todd Henderson, Paying CEOs in Bankruptcy: Executive Compensation when Agency Costs Are Low (September 2006)
307. Michael Abramowicz and M. Todd Henderson, Prediction Markets for Corporate Governance (September 2006)
308. Randal C. Picker, Who Should Regulate Entry into IPTV and Municipal Wireless? (September 2006)
309. Eric A. Posner and Adrian Vermeule, The Credible Executive (September 2006)
310. David Gilo and Ariel Porat, The Unconventional Uses of Transaction Costs (October 2006)
311. Randal C. Picker, Review of Hovenkamp, The Antitrust Enterprise: Principle and Execution (October 2006)
312. Dennis W. Carlton and Randal C. Picker, Antitrust and Regulation (October 2006)
313. Robert Cooter and Ariel Porat, Liability Externalities and Mandatory Choices: Should Doctors Pay Less? (November 2006)
314. Adam B. Cox and Eric A. Posner, The Second-Order Structure of Immigration Law (November 2006)
315. Lior J. Strahilevitz, Wealth without Markets? (November 2006)
316. Ariel Porat, Offsetting Risks (November 2006)
317. Bernard E. Harcourt and Jens Ludwig, Reefer Madness: Broken Windows Policing and Misdemeanor Marijuana Arrests in New York City, 1989–2000 (December 2006)
318. Bernard E. Harcourt, Embracing Chance: Post-Modern Meditations on Punishment (December 2006)
319. Cass R. Sunstein, Second-Order Perfectionism (December 2006)
320. William M. Landes and Richard A. Posner, The Economics of Presidential Pardons and Commutations (January 2007)
321. Cass R. Sunstein, Deliberating Groups versus Prediction Markets (or Hayek’s Challenge to Habermas) (January 2007)
322. Cass R. Sunstein, Completely Theorized Agreements in Constitutional Law (January 2007)
323. Albert H. Choi and Eric A. Posner, A Critique of the Odious Debt Doctrine (January 2007)
324. Wayne Hsiung and Cass R. Sunstein, Climate Change and Animals (January 2007)
325. Cass R. Sunstein, Cost-Benefit Analysis without Analyzing Costs or Benefits: Reasonable Accommodation, Balancing and Stigmatic Harms (January 2007)
326. Cass R. Sunstein, Willingness to Pay versus Welfare (January 2007)
327. David A. Weisbach, The Irreducible Complexity of Firm-Level Income Taxes: Theory and Doctrine in the Corporate Tax (January 2007)
328. Randal C. Picker, Of Pirates and Puffy Shirts: A Comments on “The Piracy Paradox: Innovation and Intellectual Property in Fashion Design” (January 2007)
329. Eric A. Posner, Climate Change and International Human Rights Litigation: A Critical Appraisal (January 2007)
330. Randal C. Picker, Pulling a Rabbi Out of His Hat: The Bankruptcy Magic of Dick Posner (February 2007)
331. Bernard E. Harcourt, Judge Richard Posner on Civil Liberties: Pragmatic (Libertarian) Authoritarian (February 2007)

332. Cass R. Sunstein, If People Would Be Outraged by Their Rulings, Should Judges Care? (February 2007)
333. Eugene Kontorovich, What Standing Is Good For (March 2007)
334. Eugene Kontorovich, Inefficient Customs in International Law (March 2007)
335. Bernard E. Harcourt, From the Asylum to the Prison: Rethinking the Incarceration Revolution. Part II: State Level Analysis (March 2007)
336. Cass R. Sunstein, Due Process Traditionalism (March 2007)
337. Adam B. Cox and Thomas J. Miles, Judging the Voting Rights Act (March 2007)
338. M. Todd Henderson, Deconstructing Duff & Phelps (March 2007)
339. Douglas G. Baird and Robert K. Rasmussen, The Prime Directive (April 2007)
340. Cass R. Sunstein, Illusory Losses (May 2007)
341. Anup Malani, Valuing Laws as Local Amenities (June 2007)
342. David A. Weisbach, What Does Happiness Research Tell Us about Taxation? (June 2007)
343. David S. Abrams and Chris Rohlf, Optimal Bail and the Value of Freedom: Evidence from the Philadelphia Bail Experiment (June 2007)
344. Christopher R. Berry and Jacob E. Gersen, The Fiscal Consequences of Electoral Institutions (June 2007)
345. Matthew Adler and Eric A. Posners, Happiness Research and Cost-Benefit Analysis (July 2007)
346. Daniel Kahneman and Cass R. Sunstein, Indignation: Psychology, Politics, Law (July 2007)
347. Jacob E. Gersen and Eric A. Posner, Timing Rules and Legal Institutions (July 2007)
348. Eric A. Posner and Adrian Vermeule, Constitutional Showdowns (July 2007)
349. Lior Jacob Strahilevitz, Privacy versus Antidiscrimination (July 2007)
350. Bernard E. Harcourt, A Reader's Companion to Against Prediction: A Reply to Ariela Gross, Yoram Margalioth, and Yoav Sapir on Economic Modeling, Selective Incapacitation, Governmentality, and Race (July 2007)
351. Lior Jacob Strahilevitz, "Don't Try This at Home": Posner as Political Economist (July 2007)
352. Cass R. Sunstein, The Complex Climate Change Incentives of China and the United States (August 2007)
353. David S. Abrams and Marianne Bertrand, Do Judges Vary in Their Treatment of Race? (August 2007)
354. Eric A. Posner and Cass R. Sunstein, Climate Change Justice (August 2007)
355. David A. Weisbach, A Welfarist Approach to Disabilities (August 2007)
356. David S. Abrams, More Time, Less Crime? Estimating the Deterrent Effect of Incarceration using Sentencing Enhancements (August 2007)
357. Stephen J. Choi, G. Mitu Gulati and Eric A. Posner, Professionals or Politicians: The Uncertain Empirical Case for an Elected Rather than Appointed Judiciary (August 2007)
358. Joseph Bankman and David A. Weisbach, Consumption Taxation Is Still Superior to Income Taxation (September 2007)
359. Douglas G. Baird and M. Todd Henderson, Other People's Money (September 2007)
360. William Meadow and Cass R. Sunstein, Causation in Tort: General Populations vs. Individual Cases (September 2007)
361. Richard McAdams and Janice Nadler, Coordinating in the Shadow of the Law: Two Contextualized Tests of the Focal Point Theory of Legal Compliance (September 2007)
362. Richard McAdams, Reforming Entrapment Doctrine in *United States v. Hollingsworth* (September 2007)
363. M. Todd Henderson, From *Seriatim* to Consensus and Back Again: A Theory of Dissent (October 2007)
364. Timur Kuran and Cass R. Sunstein, Availability Cascades and Risk Regulation (October 2007)
365. David A. Weisbach, The Taxation of Carried Interests in Private Equity (October 2007)
366. Lee Anne Fennell, Homeownership 2.0 (October 2007)
367. Jonathan R. Nash and Rafael I. Pardo, An Empirical Investigation into Appellate Structure and the Perceived Quality of Appellate Review (October 2007)
368. Thomas J. Miles and Cass R. Sunstein, The Real World of Arbitrariness Review (November 2007)
369. Anup Malani, Maciej F. Boni, Abraham Wickelgren, and Ramanan Laxminarayan, Controlling Avian Influenza in Chickens (November 2007)
370. Richard H. McAdams, The Economic Costs of Inequality (November 2007)

371. Lior Jacob Strahilevitz, Reputation Nation: Law in an Era of Ubiquitous Personal Information (November 2007)
372. Thomas J. Miles and Cass R. Sunstein, The New Legal Realism (December 2007)
373. M. Todd Henderson, Everything Old Is New Again: Lessons from *Dodge V. Ford Motor Company* (December 2007)
374. Jonathan Remy Nash, Economic Efficiency versus Public Choice: The Case of Property Rights in Road Traffic (December 2007)
375. Edward L. Glaeser and Cass R. Sunstein, Extremism and Social Learning (December 2007)
376. Stephen J. Choi, G. Mitu Gulati, and Eric A. Posner, Are Judges Overpaid?: A Skeptical Response to the Judicial Salary Debate (December 2007)
377. Eric A. Posner, Does Political Bias in the Judiciary Matter?: Implications of Judicial Bias Studies for Legal and Constitutional Reform (January 2008)
378. Susan Bandes, The Heart Has Its Reasons: Examining the Strange Persistence of the American Death Penalty (January 2008)
379. Susan Bandes, After Innocence: Framing Wrongful Convictions (January 2008)
380. Jacob E. Gersen and Ann Joseph O'Connell, Deadlines in Administrative Law (January 2008)
381. Richard A. Epstein, The Property Rights Movement and Intellectual Property (January 2008)
382. Richard A. Epstein, Some Reflections on Custom in the IP Universe (January 2008)
383. Richard A. Epstein, Decentralized Responses to Good Fortune and Bad Luck (January 2008)
384. Richard A. Epstein, How to Create—or Destroy—Wealth in Real Property (January 2008)
385. Richard A. Epstein, The Human and Economic Dimensions of Altruism: The Case of Organ Transplantation (January 2008)
386. Cass R. Sunstein, Adolescent Risk-Taking and Social Meaning: A Commentary (January 2008)
387. Cass R. Sunstein and Adrian Vermeule, Conspiracy Theories (January 2008)
388. Ariel Porat, Expanding Restitution: Liability for Unrequested Benefits (January 2008)
389. Randal C. Picker, *Twombly*, *Leegin* and the Reshaping of Antitrust (February 2008)
390. Alon Harrel and Ariel Porat, Aggregating Probabilities across Offences in Criminal Law (March 2008)
391. John Bronsteen, Christopher Buccafusco, and Jonathan Masur, Hedonic Adaptation and the Settlement of Civil Lawsuits (March 2008)
392. Randal C. Picker, Fair Use v. Fair Access (March 2008)
393. Jonathan S. Masur, Process as Purpose: Administrative Procedure, Costly Screens and Examination at the Patent Office (March 2008)
394. Eric A. Posner, Human Welfare, Not Human Rights (March 2008)
395. Lee Anne Fennell, Slices and Lumps, 2008 Coase Lecture (March 2008)
396. Eric A. Posner, Fault in Contract Law (March 2008)
397. Cass R. Sunstein, Is OSHA Unconstitutional? (March 2008)
398. Randal C. Picker, Take Two: Stare Decisis in Antitrust/The Per Se Rule against Horizontal Price-Fixing (March 2008)
399. M. Todd Henderson and Anup Malani, Corporate Philanthropy and the Market for Altruism (April 2008)
400. Shyam Balganesh, Copyright Incentives and Bounded Rationality: Foreseeability as a Limiting Device
401. Cass R. Sunstein and Reid Hastie, Four Failures of Deliberating Groups (April 2008)