Journal of Institutional Economics (2015), 11: 3, 561–583
© *Millennium Economics Ltd 2014* doi:10.1017/S1744137413000398
First published online 10 March 2014

Garbage in, garbage out? Some micro sources of macro errors

CASS R. SUNSTEIN*
Harvard University, Cambridge, MA, USA

REID HASTIE
Graduate School of Business, University of Chicago, Chicago, IL, USA

Abstract. Many institutions, large or small, make their decisions through some process of deliberation. Nonetheless, deliberating institutions 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. Micro mistakes can lead to macro blunders or even catastrophes. There are four such failures; all of them have implication for large-scale institutions as well as small ones. (1) Sometimes the predeliberation errors of an institution's members are amplified, not merely propagated, as a result of deliberation. (2) Institutions fall victim to cascade effects, as the initial speakers or actors are followed by their successors, who do not disclose what they know. Non-disclosure, on the part of those successors, may be a product of either informational or reputational cascades. (3) As a result of group polarization, deliberating institutions sometimes 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 institutions, shared information often dominates or crowds out unshared information, ensuring that institutions do not learn what their members know. Informational signals and reputational pressure help to explain all four errors. The results can be harmful to numerous institutions, including large ones, and to societies as a whole. Markets are able to correct some of these problems, but cascade effects occur there as well.

Many institutions, large and small, make their decisions through deliberation. When private companies or public officials are deciding how to proceed, they usually talk the problem through. And indeed, those discussions can have significant consequences for institutions – how they organize themselves and what they do. But why, exactly, is deliberation important or even desirable? When, if at all, does it correct the kinds of mistakes that are made by individuals, thinking and acting on their own?

When institutions make decisions via deliberation, they must believe that the result will be wiser judgments and better outcomes. But does deliberation

*Email: csunstein@law.harvard.edu

https://doi.org/10.1017/S1744137413000398 Published online by Cambridge University Press

actually have this effect? The idea of 'groupthink', elaborated and popularized by Irving Janis, suggests that groups may tend toward uniformity and censorship, thus failing to combine information and enlarge the range of arguments (Janis, 1983). Even markets can be vulnerable to some of the underlying problems (Shiller, 2006), as when misconceptions or rumors spread quickly and create serious problems, potentially even bubbles, at least in the short-run.

One of our principal goals here is to identify the micro sources of macro errors (compare Schelling, 2006). That goal requires an analysis that goes far beyond the idea of groupthink, one that provides a new understanding of the mechanisms that account for the failures of deliberating institutions, and that connects those failures to failures of larger institutions. Our central claim is that when institutions perform poorly, it is often because individual errors, of the kind emphasized by cognitive psychologists and behavioral economists, are replicated or even aggravated at the macro level. What we aim to establish here is that such errors, often a product of biases and heuristics, are responsible for large-scale mistakes, especially if relevant institutions do not take steps to correct them. To take just one example, the availability heuristic, used by individuals, is also used by groups, and social influences can increase its impact, even in market settings (Shiller, 2006).

We isolate two sets of influences on members of deliberating institutions. 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 and institutions 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.

We show that because of these problems, institutions often fail to achieve their minimal goal of aggregating the information actually held by those who compose them. A confident, cohesive, but error-prone group, giving effect to the mistakes of individual members, 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. And when institutions fail, one of these four problems may well be responsible. A national legislature, for example, may fall victim to cascade effects. An important government agency may not be alert to unshared information. Bubbles may be created through similar mechanisms (Shiller, 2006). There is

https://doi.org/10.1017/S1744137413000398 Published online by Cambridge University Press

a clear relationship between micro mistakes, often at the individual level, and mistakes from institutions and groups both large and small.

Successful institutions are alert, even if implicitly, to the underlying risks, and take steps to reduce them. There is a great deal of room for future empirical work, exploring how different institutions (firms, public institutions, markets of various sorts) deal with the risk of failure, perhaps through enlisting incentives, perhaps through nudges (Tahler and Sunstein, 2006), perhaps through institutional reform (Shiller, 2004). One of our goals is to suggest the importance and feasibility of such work.

1. Prospects

If deliberation is helpful within institutions, we can imagine three principal reasons:

- *The best solution wins: 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 10 people are trying to make some prediction about the future, and that one of the 10 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.
  Institutions with many members may do especially well on this count; the Internet provides many examples of ingenious 'crowd-sourcing mechanisms' that select the best solutions from large samples of 'solvers'. A good example is provided by problem-solving tournaments and bounty pay-for-solutions services, like the Netflix Prize (http://www.netflixprize.com/) and the iCisco innovation tournament (Terwiesch and Ulrich, 2009). A monetary incentive or prize structure is proposed and free-lance 'solvers' are invited to compete to produce the best solution or to exceed a criterion of achievement. These methods are revolutionizing the innovation process in many already productive firms and institutions.
- *The whole is the sum of the parts: aggregating information.* Deliberation could help to aggregate existing information in a way that leads the institution as a whole to know more than any individual does. Suppose that the institution contains no global experts on the question at issue, and that relevant information is dispersed among members so that the overarching institution is potentially expert even if none of its individual members is not. Or suppose that the institution contains a number of experts, but that each member is

https://doi.org/10.1017/S1744137413000398 Published online by Cambridge University Press

puzzled about how to solve a particular problem. Deliberation might elicit the relevant information and allow the institution 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 individual 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.

Of course a standard justification of markets, elaborated by Hayek, is that they do exceptionally well in aggregating diverse information, and hence 'know' far more than the individuals who participate in them. If markets perform better than even expert and well-motivated planners, their ability to aggregate information is an important reason. At the same time, there seem to be many information-pooling problems that are likely to exceed market capacities. And much remains unknown about the information-pooling properties of different market mechanisms. Obviously, a pari-mutuel betting market disincentivizes information-pooling, while double-auction markets (like most stock exchanges) promote sharing information, at least biased pooling to drive up the prices of options in which a person has invested.

- *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 institution 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 an institution 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. These processes are exemplified by successful brainstorming and innovation methods used by design teams at firms like IDEO. Large aggregations, again including efforts to obtain dispersed information through the Internet, also hold out considerable promise in this regard.

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* (Heath and Gonzalez, 1995). 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 (Baron *et al.*, 1996). Second, deliberation *usually reduces variance* within groups (although remember that deliberation can increase variance between groups though polarization and related effects). After talking together, group members tend to come into accord with one another.

It follows that members of deliberating institutions will usually converge on a position on which members have a great deal of confidence. This is productive 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

https://doi.org/10.1017/S1744137413000398 Published online by Cambridge University Press

turns out to be wrong (a most unfortunate and sometimes quite destructive situation). While the relevant studies typically involve members of small groups, there is every reason to believe that the two findings would hold within larger institutions, at least if their decisions are a product of some form of deliberation.

Unfortunately, there is no systematic evidence that deliberating groups or institutions 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 (Gigone and Hastie, 1997). Hence, it is false to say that group members usually end up deferring to their internal experts. Truth does not always 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' (MacCoun, 2002: 116, 121) (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.

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 (MacCoun, 2002). In this respect, there is a risk that institutions will not do better than their individual members, and that micro errors will produce macro errors as well. 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' (Armstrong, 2001: 433). There is a lesson here for the operation of large institutions, including businesses and government bodies.

2. 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

https://doi.org/10.1017/S1744137413000398 Published online by Cambridge University Press

on the ground that their own judgments must be ill-informed of wrong. Large institutions, including nations, might be vulnerable in this way.

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. 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.

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 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.

In public institutions, especially but not only in authoritarian nations, there is a risk that the problem of self-silencing can be especially serious. If officials do not tell other officials what they know, leaders may not end up with the information that they need in order to decide well. One of the advantages of markets, in this regard, is that those with private information are free and incentivized to choose in a way that reflects what they know. Self-silencing is unlikely to be problem when people are making investment decisions, in which self-interest strongly inclines people to act on the basis of the information that they have.

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 (Christenson and Abbott, 2000). 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. Of course this can be a problem for large institutions as well as small groups. Consider the finding that one of the most common defensive tactics for small coalition or solo

https://doi.org/10.1017/S1744137413000398 Published online by Cambridge University Press

hold-out jurors is to refuse to discuss the case further with other jurors (Hastie *et al.*, 2000).

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

3. 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 subject to identifiable biases, which also produce errors (Gilovich *et al.*, 2002). In putting a spotlight on these biases, we mean to point to some of the micro foundations of macro errors. The point is analogous to that made by Schelling (2006) and his followers (Watts, 2011), who emphasize how individually rational decisions can produce surprises and mistakes. What we mean to add is that institutions and groups may not correct bounded rationality at the individual level (Shiller, 2006). As we shall see, there are important interactions between the work of Schelling and his followers and various findings about departures from rationality (Watts, 2011).

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 (Tversky and Kahneman, 1973). The point bears on private and public responses to 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 (Slovic, 2000). Similarly, earlier events will have a smaller impact than later ones. The point helps to explain both private and institutional behavior. For example, whether people will buy insurance for natural disasters is greatly affected by recent experiences (Slovic, 2000). 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 fade. 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 not unpredictable that before that disaster, such steps would be quite inadequate. Here, then, we can see that the availability heuristic

https://doi.org/10.1017/S1744137413000398 Published online by Cambridge University Press

operates at the level of large institutions. In deciding whether to take steps to combat a potentially serious national problem, the availability heuristic may have a significant impact.

Most people also follow the representativeness heuristic, in accordance with which judgments of probability are influenced by assessments of similarity (the extent to which A 'looks like' B) (Tversky and Kahneman, 1982). 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 (Rozin and Nemeroff, 2002). The representativeness and availability heuristics often work well, but they are certainly not optimal and they can lead to severe blunders.

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 often they do not – a vivid illustration of the principle, 'garbage in, garbage out', in a way that mocks the aspiration to collective or institutional correction of individual blunders. In fact, individual errors are not merely replicated but actually *amplified* in many group or institutional decisions – a process of 'some garbage in, much garbage out'. We can see this problem in multiple domains.

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 (MacCoun, 2002). Groups have been found to amplify reliance on the representativeness heuristic (Stasson *et al.*, 1988); to exhibit even larger framing effects than individuals (Kerr *et al.*, 1996); to show more overconfidence than group members (Sniezek and Henry, 1989) (Johnson, 2009); to be more affected by the biasing effect of spurious arguments from lawyers (Schumann and Thompson, 1989); to be more susceptible to the 'sunk cost fallacy' (Whyte, 1993); and to produce more preference reversals in choice (Irwin and Davis, 1995; Mowen and Gentry, 1980). 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 (Whyte, 1986). In addition, and importantly, groups demonstrate essentially the same level of reliance on the availability heuristic, even when use of that heuristic leads to clear errors (Stasson *et al.*, 1988) – a point that has implications for markets (Shiller, 2006).

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 even more so, if members identify strongly with the groups to which they belong* (Stasser and Dietz-Uhler, 2001). 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

https://doi.org/10.1017/S1744137413000398 Published online by Cambridge University Press

on its misguided course simply because of group dynamics. (Enron is an obvious 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 level of groups and institutions? 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 or proper. 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, these factors are part of the explanation.

To be sure, deliberating groups can attenuate certain biases, and this evidence suggests circumstances in which well-designed institutions can perform well. We have seen that for eureka problems, groups often find the right answer, even if most individual members begin with a biased answer. 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 for the answer 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 benefit of hindsight, that they would have accurately predicted the outcome of an event (a political election, a natural disaster, an illness, a change in the stock market). Compared to individuals, groups are less susceptible to hindsight bias (Stahlberg *et al.*, 1995). Apparently group members who are less susceptible to that bias are able to persuade others that it is indeed a bias. It is possible that large institutions, assessing past events, will be less likely to fall prey to hindsight bias.

But the broader point is that with group discussion, individual errors are often 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.

4. 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

https://doi.org/10.1017/S1744137413000398 Published online by Cambridge University Press

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. Both of these processes affect private and public institutions.

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 (Hirshleifer, 1995).

To see how informational cascades work, imagine an institution that is deciding whether to authorize some new venture (Hirshleifer, 1995). Let us also assume that the institution's 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 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 (Hirshleifer, 1995).

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 institution – 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 the institution's members know that it is a terrible idea. The reason is that

https://doi.org/10.1017/S1744137413000398 Published online by Cambridge University Press

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 member of the institution, 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 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. The most diabolical situations occur when information is mostly unshared and where pooling information is necessary to achieve an effective solution. Under these, so-called 'hidden profile' situations, every individual can use his or her personal information rationally and still reach an incorrect individual decision (see further discussion of 'hidden profiles' below). Then the negative consequences are amplified, sometimes to a catastrophic level, by the social cascade process. Cascades often occur in the real world within public and private institutions and deliberating groups (Hirshleifer, 1995; Sunstein, 2003); they can also affect the movements of stock markets, at least temporarily (Shiller, 2006).

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 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.

Suppose Albert suggests that climate change 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 climate change 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 to lose their positive regard.

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 institutional 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.

https://doi.org/10.1017/S1744137413000398 Published online by Cambridge University Press

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 with macro consequences (Kuran and Sunstin, 1999).

In the area of risk, availability cascades are common, and they affect important social institutions, including national legislatures. 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. Large-scale institutions can move as a result. In business, availability cascades are also common. A particular event, a dramatic success, or a singular failure may spread like wildfire within or across firms, leading to a judgment about other apparently similar events or products. If a movie (*Twilight*), 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 to fruition, but the instance may not be representative. 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 a particular socially themed situation comedy has succeeded, another one on the same theme 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. Such cascades can also be found in market settings, given rise to bubbles (Shiller, 2006).

5. 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*

https://doi.org/10.1017/S1744137413000398 Published online by Cambridge University Press

tended before deliberation began (Brown, 1986). The problem is especially severe for groups of like-minded people, who typically end up in more extreme positions as a result of deliberation. And while most of the relevant studies involve small groups, there is no question that group polarization occurs within large institutions.

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 (Brown, 1986). 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 (Brown, 1986).

The original psychological experiments on the effects of deliberation are relevant to businesses and governments alike. They involved risk-taking behavior, with a demonstration that risk-inclined people become still more risk-inclined after they deliberate with one another (Brown, 1986). 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 threw 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 (Moscovici and Zavalloni, 1969).

A straightforward interpretation turned out to be able to reconcile these competing findings: the predeliberation median is the best predictor of the direction of the shift (Brown, 1986). When more than half of 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 predeliberation 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

https://doi.org/10.1017/S1744137413000398 Published online by Cambridge University Press

attack will occur in the United States in the next year, group polarization will not be easy to test, simply because the answer is either yes or no, and it is not easy 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 (Brown, 1986).

Note here that in the United States, an important institution has been found to be vulnerable to group polarization. Federal judges are highly susceptible to such 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 (Sunstein, 2006). Juries polarize as well (Schkade *et al.*, 2000).

Why does group polarization occur? There are three reasons (Brown, 1986; Glaeser and Sunstein, 2009). 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 and even old arguments, when repeated, increase belief in what is asserted (Dechêne *et al.*, 2010). 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. This problem can of course occur within large institutions as well as small groups.

The second explanation is reputational and 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, slightly because they want to preserve their self-presentation. But they shift accordingly (Brown, 1986).[1]

1 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 (Baron *et al.*, 1996).

https://doi.org/10.1017/S1744137413000398 Published online by Cambridge University Press

The third explanation of group polarization stresses the close links among confidence, extremism, and corroboration by others (Baron *et al.*, 1996). 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 (Baron *et al.*, 1996). Many minds can badly blunder in this way, including the many minds who play a role in political choices.

Does group polarization lead to accurate or inaccurate answers? Do institutions err when they polarize? No general answer would make sense. Everything depends on the relationship between the correct answer and the institution's predeliberation tendencies. If the institution is leaning toward the right answer, polarization might lead its members 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, including by large institutions.

6. Deliberative failure 4: over-weighting common knowledge

The last failure is in many respects the most interesting. Suppose that people within an institution 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 institutions will not perform well if they focus on shared information and neglect information that is held by one or a few people. Unfortunately, countless studies demonstrate that this regrettable result is highly likely (Stasser and Titus, 2003).

'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 (Gigone and Hastie, 1993; Stasser and Titus, 2003). The most obvious explanation of the effect is the simple fact that as a statistical matter, common knowledge is more likely to be communicated to the group; but social influences play a role as well.

Hidden profiles

Consider a study of serious errors within working groups, both face-to-face and online (Hightower and Sayeed, 1995). The purpose of the study was to

https://doi.org/10.1017/S1744137413000398 Published online by Cambridge University Press

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 showed 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. But the distribution of the information among members was rigged by the experimenters so that each member had *a personal sample of attributes* that favored choosing an inferior candidate, although the full set of all information *across all members* strongly indicated the correct candidate. The groups consisted of three people, some operating face-to-face, sometimes 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 acomplications and fuel debate' (Wallace, 2001: 82).

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 (Stasser and Titus, 1985). In the first 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 (Stasser and Titus, 1985, 2003). This is a clear example of appropriate aggregation of information. Groups significantly outperformed individuals, apparently because of the exchange of information and reasons. Here, then, is a clear illustration of the possibility that groups can aggregate what members know in a way that produces sensible outcomes.

In the second condition, by contrast, the information that favored candidate A was parceled out to various members of the group so that only 33% 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% 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 (Stasser and Titus, 1985). 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

https://doi.org/10.1017/S1744137413000398 Published online by Cambridge University Press

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' (Stasser and Titus, 1985). 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' (Stasser and Titus, 1985). 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 (Stasser and Titus, 1985).

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' (Stasser and Titus, 2003). Here too, the evidence involves small groups. But there is every reason to believe that large institutions are vulnerable to the same problem.

The common-knowledge effect

The foregoing results are best understood as a consequence of the common-knowledge effect, by which information held by all group members has far more influence on group judgments than information held by one member or a few (Gigone and Hastie, 1993). 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' (Gigone and Hastie, 1993: 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' (Gigone and Hastie, 1993: 973).

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 (Gigone and Hastie, 1993). 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 and within private and public institutions? Affirmative evidence comes from a hiring exercise involving high-level executives (Abele *et al.*, 2005). In this

https://doi.org/10.1017/S1744137413000398 Published online by Cambridge University Press

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. 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', (Kameda *et al.*, 1997) in the sense that their knowledge is shared 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 participate at low levels.

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 (Stasser *et al.*, 1989). For this reason, groups can go very badly wrong, especially 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' (Stasser *et al.*, 1989: 78).

It follows that for very large groups and even institutions – 12 people, 30 people, and 200 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 errors are a bad combination.

https://doi.org/10.1017/S1744137413000398 Published online by Cambridge University Press

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 the ultimate decision, while initially unshared information will have disproportionately low impact. 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' (Stasser *et al.*, 1989: 78). 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 business and government, 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 reluctant to introduce 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 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.

The other side of this coin is that group members will be positively motivated to mention and discuss information that is already shared. After all, this is information on which they can compare interpretations and implications. Thus, shared information will further crowd out unshared information as the group analyzes and reanalyzes what its members already all know, magnifying the

https://doi.org/10.1017/S1744137413000398 Published online by Cambridge University Press

effect of that knowledge and blocking the introduction of new, potentially more informative facts.

This point creates a big problem for institutions, which will fail to get information that they need, and which can therefore lead to error. Consider the finding that group members typically *underestimate* the performance of low-status and low-participating 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 (Ridgeway, 2001).

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 (Wittenbaum *et al.*, 1999). 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 communications, people gave 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' (Stasser and Titus, 2003: 311). The general problem is that deliberating groups often perform poorly because they fail to elicit information that could steer them in the right directions.

7. Conclusion

Our goal here has been to provide a novel understanding of the relationship between micro errors and institutional blunders. Individual errors, of the kind emphasized by behavioral economists, have important implications for institutions, both large and small.

We have identified four common failures. (1) Sometimes individual errors are amplified, not merely propagated, as a result of deliberation. (2) Institutions fall victim to cascade effects, as the early speakers or actors prevent others from learning what is known by their successors. (3) Because of group polarization, members of deliberating institutions often end up in a more extreme position in line with their predeliberation tendencies. (4) In deliberation, shared information usually dominates or crowds out unshared information, ensuring that institutions do not learn what their members know.

In the future, there is a great deal of room for empirical work, exploring the role of these problems in different institutions, including firms, governmental bodies, and markets themselves (Shiller, 2006). Illuminating comparisons could undoubtedly be shown. Bubbles of multiple kinds arise through several of the mechanisms we have elaborated (ibid.).

As we have seen, informational signals and reputational pressure help to explain all four errors. A disturbing result is that institutions often do not improve on, and sometimes do worse than, their average or median member.

https://doi.org/10.1017/S1744137413000398 Published online by Cambridge University Press

An important task, for many deliberating groups, is to devise strategies that will overcome the relevant problems and increase the likelihood that institutional judgments will take advantage of the knowledge held by all group members. We have not said anything here about how that task might be accomplished. Our goal has been instead to specify links between micro errors and micro failures, and to understand the mechanisms that threaten to make the resulting problems so challenging to solve.

Acknowledgements

We are grateful to Teppo Felin and two anonymous reviewers for valuable comments on a previous draft. We draw at some points on the discussion in Cass R. Sunstein, *Infotopia* (New York: Oxford University Press, 2006).

References

Abele, S., G. Stasser, and S. Vaughan-Parsons (2005), 'Information Sharing, Cognitive Centrality, and Influence among Business Executives during Collective Choice', http://www.erim.eur.nl (Accessed 9 December, 2013).

Armstrong, J. (2001), 'Combining Forecasts', in J. Armstrong (ed.), *Principles of Forecasting: A Handbook for Researchers and Practitioners*, Boston: Kluwer Academic.

Baron, R. S., S. I. Hoppe, C. F. Kao, B. Brunsman, B. Linneweh, and D. Rogers (1996), 'Social Corroboration and Opinion Extremity', *Journal of Experimental Social Psychology*, **32**(6): 537–560.

Brown, R. (1986), *Social Psychology: The Second Edition*, New York: Free Press.

Christenson, C. and A. Abbott (2000), 'Team Medical Decision Making', in G. Chapman and F. Sonnenberg (eds.), *Decision Making in Health Care*, New York: Cambridge University Press.

Dechêne, A., C. Stahl, J. Hansen, and M. Wänke (2010), 'The Truth About the Truth: A Meta-Analytic Review of the Truth Effect', *Personality and Social Psychology Review*, **14**(2): 238–257.

Gigone, D. and R. Hastie (1993), 'The Common Knowledge Effect: Information Sharing and Group Judgment', *Journal of Personality and Social Psychology*, **65**(5): 959.

Gigone, D. and R. Hastie (1997), 'Proper Analysis of the Accuracy of Group Judgments', *Psychological Bulletin*, **121**(1): 149.

Gilovich, T., D. Griffin, and D. Kahneman (eds.) (2002), *Heuristics and Biases: The Psychology of Intuitive Judgment*, Cambridge, UK: Cambridge University Press.

Glaeser, E. L. and C. R. Sunstein (2009), 'Extremism and Social Learning', *Journal of Legal Analysis*, **1**(1): 263–324.

Hastie, R. (1986), 'Experimental Evidence of Group Accuracy', in B. Grofman and G. Owen (eds.), *Information Pooling and Group Decision Making*, Greenwich, CT: JAI Press, p. 129.

Heath, C. and R. Gonzalez (1995), '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**(3): 305–326.

https://doi.org/10.1017/S1744137413000398 Published online by Cambridge University Press

Hightower, R. and L. Sayeed (1995), 'The Impact of Computer-Mediated Communication Systems on Biased Group Discussion', *Computers in Human Behavior*, **11**(1): 33–44.

Hirshleifer, D. (1995), 'The Blind Leading the Blind: Social Influence, Fads and Informational Cascades', in M. Tommasi and K. Ierulli (eds.), *The New Economics of Human Behaviour*, Cambridge, UK: Cambridge University Press.

Irwin, J. R. and J. H. Davis (1995), 'Choice/Matching Preference Reversals in Groups: Consensus Processes and Justification-Based Reasoning', *Organizational Behavior and Human Decision Processes*, **64**(3): 325–339.

Janis, I. L. (1983), *Groupthink*, Boston: Houghton Mifflin, pp. 2–13.

Johnson, D. D. (2009), *Overconfidence and War: The Havoc and Glory of Positive Illusions*, Cambridge: Harvard University Press.

Kameda, T., Y. Ohtsubo, and M. Takezawa (1997), 'Centrality in Sociocognitive Networks and Social Influence: An Illustration in a Group Decision-Making Context', *Journal of Personality and Social Psychology*, **73**(2): 296.

Kerr, N. L., R. J. MacCoun, and G. P. Kramer (1996), 'Bias in Judgment: Comparing Individuals and Groups', *Psychological Review*, **103**(4): 687.

Kuran, T. and C. R. Sunstein (1999), 'Availability Cascades and Risk Regulation', *Stanford Law Review*, **51**(4): 683–768.

MacCoun, R. J. (2002), 'Comparing Micro and Macro Rationality', in R. Gowda and J. Fox (eds.), *Judgments, Decisions, and Public Policy*, Cambridge, UK: Cambridge University Press.

Moscovici, S. and M. Zavalloni (1969), 'The Group as a Polarizer of Attitudes', *Journal of Personality and Social Psychology*, **12**(2): 125.

Mowen, J. C. and J. W. Gentry (1980), 'Investigation of the Preference-Reversal Phenomenon in a New Product Introduction Task', *Journal of Applied Psychology*, **65**(6): 715.

Ridgeway, C. L. (2001), 'Social Status and Group Structure', *Blackwell Handbook of Social Psychology: Group Processes*, 352–375. Available at: http://onlinelibrary.wiley.com/doi/10.1002/9780470998458.ch15/summary

Rozin, P. and C. Nemeroff (2002), 'Sympathetic Magical Thinking: The Contagion and Similarity "Heuristics"', in T. Gilovich, D. Griffin and A. Kahneman (eds.), *Heuristics and Biases: The Psychology of Intuitive Judgment*, Cambridge, UK: Cambridge University Press.

Schkade, D., C. R. Sunstein, and D. Kahneman (2000), 'Deliberating About Dollars: The Severity Shift', *Columbia Law Review*, **100**(4): 1139–1175.

Schelling, T. (2006), *Micromotives and Macrobehavior*, New York: W.W. Norton.

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

Shiller, R. (2004), *The New Financial Order*, New Haven: Yale University Press.

Shiller, R. (2006), *Irrational Exuberance*, New York: Crown Business.

Slovic, P. E. (2000), *The Perception of Risk*, New York: Earthscan Publications.

Sniezek, J. A. and R. A. Henry (1989), 'Accuracy and Confidence in Group Judgment', *Organizational Behavior and Human Decision Processes*, **43**(1): 1–28.

Stahlberg, D., F. Eller, A. Maass, and D. Frey (1995), 'We Knew it All Along: Hindsight Bias in Groups', *Organizational Behavior and Human Decision Processes*, **63**(1): 46–58.

Stasser, G. and B. Dietz-Uhler (2001), 'Collective Choice, Judgment, and Problem Solving', *Blackwell Handbook of Social Psychology: Group Processes*, **3**: 31–55.

https://doi.org/10.1017/S1744137413000398 Published online by Cambridge University Press

Stasser, G., L. A. Taylor, and C. Hanna (1989), 'Information Sampling in Structured and Unstructured Discussions of Three-And Six-Person Groups', *Journal of Personality and Social Psychology*, **57**(1): 67.

Stasser, G. and W. Titus (1985), 'Pooling of Unshared Information in Group Decision Making: Biased Information Sampling During Discussion', *Journal of Personality and Social Psychology*, **48**(6): 1467.

Stasser, G. and W. Titus (2003), 'Hidden Profiles: A Brief History', *Psychological Inquiry*, **14**(3–4): 304–313.

Stasson, M. F., K. Ono, S. K. Zimmerman, and J. H. Davis (1988), 'Group Consensus Processes on Cognitive Bias Tasks: A Social Decision Scheme Approach', *Japanese Psychological Research*, **30**(2): 68.

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

Sunstein, C. R. (ed.) (2006), *Are Judges Political?[Electronic Resource]: An Empirical Analysis of the Federal Judiciary*, Washington, D. C.: Brookings Institution Press.

Tahler, R. and C. R. Sunstein (2008), *Nudge: Improving Decisions About Health, Wealth, and Happiness*, New Haven: Yale University Press.

Terwiesch, C. and K. T. Ulrich (2009), *Innovation Tournaments: Creating and Selecting Exceptional Opportunities*, Cambridge: Harvard Business Press.

Tversky, A. and D. Kahneman (1973), 'Availability: A Heuristic For Judging Frequency and Probability', *Cognitive Psychology*, **5**(2): 207–232.

Tversky, A. and D. Kahneman (1982), 'Judgment Under Uncertainty: Heuristics and Biase', in D. Kahneman, P. Slovic, and A. Tversky (eds.), *Judgment Under Uncertainty: Heuristics and Biases*, Cambridge, UK: Cambridge University Press.

Wallace, P. (2001), *The Psychology of the Internet*, Cambridge, UK: Cambridge University Press.

Watts, D. (2011), *Everything Is Obvious*, New York: Crown Business.

Whyte, G. (1986), 'Escalating commitment to a course of action: A reinterpretation', *Academy of Management Review*, **11**(2): 311–321.

Whyte, G. (1993), 'Escalating Commitment in Individual and Group Decision Making: A Prospect Theory Approach', *Organizational Behavior and Human Decision Processes*, **54**(3): 430–455.

Wittenbaum, G. M., A. P. Hubbell, and C. Zuckerman (1999), 'Mutual Enhancement: Toward an Understanding of the Collective Preference for Shared Information', *Journal of Personality and Social Psychology*, 77(5): 967.

https://doi.org/10.1017/S1744137413000398 Published online by Cambridge University Press