METAPHILOSOPHY

Published by Blackwell Publishing Ltd, 9600 Garsington Road, Oxford OX4 2DQ, UK, and 350 Main Street, Malden, MA 02148, USA
METAPHILOSOPHY
Vol. 44, No. 3, April 2013
0026-1068

INFOSTORMS

PELLE G. HANSEN, VINCENT F. HENDRICKS, AND RASMUS K. RENDSVIG

Abstract: It has become a truism that we live in so-called information societies where new information technologies have made information abundant. At the same time, information science has made us aware of many phenomena tied to the way we process information. This article explores a series of socio-epistemic information phenomena resulting from processes that track truth imperfectly: pluralistic ignorance, informational cascades, and belief polarization. It then couples these phenomena with the hypothesis that modern information technologies may lead to their amplification so as to give rise to what are called "infostorms." This points to the need for studying further the exact relations between information technologies and such infostorms, as well as the ways we may design technologies to avoid being misled away from what we have good reasons to believe.

Keywords: information, pluralistic ignorance, informational cascades, belief polarization.

1. Introduction

1.1. Three Examples

In 1995, a book of little note—*The Discipline of Market Leaders: Choose Your Costumers, Narrow Your Focus, Dominate Your Market*, written by two market gurus, Michael Tracy and Fred Wiersema—suddenly came in at no. 8 on the *New York Times* best-seller list, where it remained for fifteen weeks, and at no. 1 on the *Businessweek* best-seller list. This was despite the fact that several reviews had rated it as mediocre or even bad.

In 2007, the long forgotten and also unnoteworthy book entitled *Love Letters of Great Men and Women: From the Eighteenth Century to the Present Day*, collected in the 1920s by C. H. Charles, suddenly climbed Amazon.com's best-seller list, ultimately peaking in its rise to the stars at no. 134. The book was bought by thousands of people and quickly sold out. The interesting thing about the book was not its content but that it was bought by accident by thousands of consumers looking for another book that didn't exist.

In April 2011, a postdoc student at U.C. Berkeley logged on to Amazon to buy his lab an extra copy of yet a third book of little note, at least to anyone working outside *Drosophila* developmental biology, *The Making of a Fly: The Genetics of Animal Design* (1992), by Peter A. Lawrence. Although a classic work in developmental biology and out of print, two retailers could offer a new copy, but at prices that were surprisingly high: $1,730,045.91 and $2,198,177.95, respectively (plus $3.99 for shipping). Even more surprising was that the prices at which the book was offered increased even more for each day that passed. On April 18, the price ultimately peaked when Bordeebook offered the book at a startling $23,698,655.93.

1.2. Social Information Processes

What connects these three events is not only that they all concern books. What is more important is that they concern how *social information processes* may affect individual beliefs in large groups of agents in ways that track truth imperfectly, so that people end up believing false propositions and as a result possibly act contrary to their goals and interests when faced with uncertainty.

In the first case, agents in the book market came to believe that the book *The Discipline of Market Leaders* was worth reading because it featured on the otherwise highly credible best-seller lists of the *New York Times* and *Businessweek*. What they did not know was that the two authors, Michael Tracy and Fred Wiersema, had actually bought ten thousand copies of their own book at the bookstores from whose sales these best-seller lists were compiled—lists that are compiled in order to direct the buying behavior of agents in an inherently uncertain market by providing information about what other readers have bought. As a result, agents in the book market in search of a good and useful book on marketing came to believe that the book was worth reading and hence buying—according to most reviewers, a false belief that was only corrected when it was too late and when their buying behavior had already fed into the best-seller system itself.

The best-seller list is just one technology or information system for handling a massive book market and shortcutting the tiresome process of reading book reviews. It simply samples bookstores' sales to determine which books are selling the best, with the aim of advising people as to what is a good buy. To work as intended, however, a system like the best-seller list depends crucially on its ability to reflect or lead us to the "truth" (which books are worth buying) on the basis of its input (books sold). Thus, the system works like a social heuristic for individual decision making by trying to answer a complex question with a simpler one. The authors of *The Discipline of Market Leaders* knew how to shortcut this system to their own advantage and to the disadvantage of the users of the best-seller lists.

The second case, of *Love Letters of Great Men and Women*, resulted from the automatic pairing by computers of offers on Amazon triggered by a scene in the 2007 movie *Sex and the City* and based on search words by users. In the scene, the main character, Carrie Bradshaw, reads a book entitled *Love Letters of Great Men*, from which her husband-to-be, "Mr. Big," later sends her quotes by e-mail. As a result, thousands of fans of the movie logged on to Amazon and searched for the book to get a piece of the big city romance. Unfortunately, no such book turned up in their search, as the book actually did not exist. Instead the search engine on Amazon suggested the 1920s collection *Love Letters of Great Men and Women* as a possible match, leading multiple customers to click on this entry and ultimately some of them to buy the antiquated version. This in turn lead Amazon's computers to automatically pair the book in special offers with various merchandise related to the *Sex and the City* franchise, leading even more customers to believe that this was actually the book they were in search of and hence to buy it. Soon the book rose to a prominent place on Amazon's best-seller list as the no. 134 best-selling book.

1.3. The Concept of Infostorms

The hypothesis is that modern information technologies have magnified and amplified phenomena for which social information processes threaten to distort truth, making us more vulnerable to err than ever, and on a much larger scale. The abundance of information driven by technologies such as the fast printing press, radio, computers, and in particular the World Wide Web has forced us to increasingly rely on information technologies that shortcut traditional cumbersome search processes unable to cope with the abundance of available information, and that offer tempting avenues for bypassing traditional slow gatekeepers of truth. Relying more and more on social information technologies or systems like these not only makes such sidetracking possible and more likely to occur, it also increases the numerical reach, if not the proportions, of the spreading of false beliefs and consequences thereof, intentional or nonintentional. When this happens we call the resulting phenomena "infostorms."

The problem is that while the information phenomena magnified by such technologies have always existed, they now take on new proportions, with possibly severe consequences for the democratic institutions underpinning the information societies we live in. The more we uncritically rely on automatic information technologies, the more likely it is that the consequences go unnoticed, sometimes with absurd results.

This uncritical reliance is what produced the third case above, when Peter Lawrence's *The Making of a Fly* reached the startling price of $23,698,655.93. Behind this absurd price setting was the use of automatic price-setting algorithms by two retailers—Bordeebook and Profnath—

who had set their prices on the book conditional upon each other by 0.9983 and 1.270589, respectively. This automatization of price adjustment led to the gradual increase in price that ultimately resulted in the absurd valuation.

While this example, like the other two, did not have severe consequences for our democratic institutions, it exemplifies what may happen to the reflection of truth—*in casu*, the true market value—when we give our decision-making power to, and rely unconditionally on, information technologies and processes. In other cases the result of relying on such processes may amplify information phenomena that track truth imperfectly in ways that give us reasons to believe the truly unbelievable, with severe consequences for society.

This article explores three such information phenomena. It then couples them with the hypothesis that modern information technologies may give rise to their amplification so as to give rise to infostorms. If this empirical hypothesis should turn out to be true, it follows that the social epistemological study of information phenomena and the way they interact with modern information technologies becomes one of the most important research areas of our time. Hence, analogously with the biases and heuristics program of behavioral economics, which has turned out to offer insights of individual decision making important on a societal level, the infostorms program potentially offers insights on social decision making with relevancy on the same scale, but even for areas where agents act rationally. Ultimately this leads us to conclude that there is an urgent need to study further the exact relations between information phenomena, information technologies, and infostorms, as well as the ways we may design these technologies to avoid being misleading. As will be evident from our discussion, we believe that formal epistemology has a central role to play in this effort.

2. Information Phenomena

2.1. Rational Interaction and Information Phenomena

Rational action is determined by the knowledge agents have, their preferences, and the arguments they can muster for their opinions, decisions, and actions. However, societies—modern as well as traditional—are based on the fact that individual decision making unavoidably takes place in social settings comprising the interactions among agents in various structures. Thus, communication and intelligent information processing are prerequisites for informed decision making, carrying out important actions, and obtaining true beliefs—all of which are cornerstones of rational human interaction.

Yet, recent studies in social psychology, social science, economics, computer science, and jurisprudence show that concepts central to rational

collective behavior since the Enlightenment, like qualified decision, informed action, truthful justification, and so on, are acutely sensitive to the way in which agents or members of a group *process* their information in order to rationally interact.

The notion of an information phenomenon covers robust and reproducible phenomena of belief configurations and dynamics that result from the flow, exchange, and interaction in social processes and systems of information, reasonable beliefs, and their consequent actions. In particular, we will focus here on such phenomena when they track truth imperfectly by giving rational agents reason to believe false propositions and thus result in actions that miss their intended purpose.

Given this definition, information phenomena are conceived differently from cognitive biases, since they do not result from the biased individual processing of information or perception. Rather, they are consequences of the rational workings of what Kahneman has labeled as system 2 thinking in social settings (Kahneman 2011).

2.2. Derailing Rational Agents

The abovementioned empirical findings, especially those from social psychology and economics, demonstrate in various ways how the beliefs of agents may become derailed from the truth. Core notions that have been produced by these efforts include *pluralistic ignorance* (Katz and Allport 1931; Krech and Crutchfield 1948; Halbesleben and Buckley 2004), *informational cascades* (Bikhchandani, Hirshleifer, and Welch 1992; Centola, Willer, and Macy 2005), and *belief polarization* (Cooper, Kimberly, and Weaver 2004; Sunstein 2009).

Following Hansen and Hendricks 2011 we may classify these various phenomena until recently studied primarily by social psychology *as components of information phenomena* according to how they are produced by information processes.

Information phenomena		***Information problem***		***By information process***
Informational cascades	generated by	too much information	coupled with	social proof
Bystander effects	generated by	too little information	coupled with	social proof plus pluralistic ignorance
Belief polarization/ extremism	generated by	information selection	coupled with	echochamber

Such phenomena, and others like them, are potentially dangerous to collective deliberation, decision making, and action, since they may, with very unfortunate collective consequences, tap into the way in which "informed" agents make "rational" decisions, perform "rational" actions, and hold "rational" beliefs. Further, if embedded in information technologies

each one of these phenomena is likely to generate infostorms. Fortunately, it seems that these "rational" pitfalls may be resolved by dealing with them as information control problems. However, in order to control information in the right way, one has to properly identify and analyze the structure of information problems as well as the information processes involved. To this end, formal epistemology is beginning to make some significant contributions.

3. The Structure of Information Problems and Processes

Each one of these phenomena comes with a structure, some of which may be properly characterized by formal means.

3.1. Informational Cascades

The notion *informational cascade* can be interpreted to cover a wide range of different phenomena. Among these are the best-seller example mentioned in the introduction, jaywalking, changing your mind about eating at an empty restaurant because the place across the street is close to filled (Banerjee 1992), and in general using popularity as a measure of quality.

All these examples are empirical in nature. However, the original definition of an informational cascade was given in relation to a specific behavior of Bayesian-rational agents in a mathematically defined setup (Bikchandani, Hirshleifer, and Welch 1992), and by now a variety of different models exist (see, e.g., Smith and Sørensen 2011; Anderson and Holt 1997; Banerjee 1992; Easley and Kleinberg 2010).

3.1.1. Structure

In general terms, the structure underlying informational cascades consists of

1. a set of rational agents that act sequentially,
2. a set of options between which the agents can choose, and
3. a preference order on the outcome of each choice.

Typically, agents are modeled as Bayesian maximizers of expected utility. The decision is made under uncertainty in the sense that no agent knows which action leads to the jointly preferred outcome. That there is a jointly preferred outcome is essential when it comes to the epistemic assumptions made. There is no strategic interaction in the decision problem, so no agents will have an incentive to mislead later agents by choosing contrary to the best of their knowledge. This in turn means that subsequent agents may base their decision not only on their private information but also on the action of those that act before them. Specifically, the following epistemic assumptions are in order:

1. the underlying structure is known to all agents; the sequence of agents is known to all,
2. each agent makes a rational decision based on available information, which consists of
   a. a *private signal* about which action will lead to which outcome, which is known to be more often right than it is wrong;
   b. a *public signal* consisting of the string of actions performed by the previous agents,
3. knowledge among the agents that their signals are equally likely to be correct, and
4. knowledge of rationality as described in 2.

Notice that in *b* it is only the *actions*, not the *signals*, of previous agents that can be observed. Notice furthermore the fact that the sequence of agents is known to all is in conjunction with *b* taken to imply that any agent knows what public signal any previous agent received.

A *run* of such a model may be conceived as a line of agents, each waiting to make a decision between a (finite) set of choices. In runs where later agents choose to ignore their private information and act on the information conveyed by previous agents' actions, an *informational cascade* is said to be in effect.

3.1.2. Illustration: Initiating a Cascade

To illustrate, let us consider a situation where the agents have to make a binary choice between turning left or turning right at a junction in a maze—or just get off an airplane. Before receiving their private signal of *left* or *right*, each agent will be indifferent to the two options. When the first agent receives her private signal, say *left*, she will take this to indicate the correct path out of the maze. Given that she has no further information available, she will follow her private signal, thereby conveying a *left* action to all subsequent agents.

When the second agent must choose, the public signal of an executed *left* action in conjunction with knowledge of rationality may be used to deduce that the first agent's signal was *left*. Two situations may now have occurred: one in which the second agent received the private signal *left*, in which case he should choose to go left, or one where he received private signal *right*, in which case his available information—a *left* signal from agent 1 and a *right* signal from himself—will suggest opposite responses. Since both signals are known by agent 2 to be equally likely to be correct, rationality specifies no concrete plan of action. Hence the agent must choose based on some *tie-breaking rule*, for example, by randomizing, choosing to follow his private signal, and so on. The epistemic assumptions regarding tie-breaking rules are discussed below. For now, assume that the second agent received a *left* signal, and therefore chooses to go left.

The actions of agents 1 and 2 send a public (*left*, *left*) signal to agent 3. Like agent 2, agent 3 can deduce the private signal of agent 1. In addition, given suitable assumptions regarding the tie-breaking rule, agent 3 may also deduce that agent 2 received a *left* signal. As it is known that every private signal is equally likely to be correct, *it now does not matter for agent 3's action what signal she received*. If agent 3 received a *left* signal, she too should choose to go left. If she received a *right* signal, the information extrapolated from the public (*left*, *left*) signal results in left still being more probable than right. She will therefore choose to ignore her private information and act in accordance with the group behavior. Thereby agent 3 will be the first agent in an informational cascade.

Upon receiving the (*left*, *left*, *left*) action string, agent 4 will also choose to ignore his private signal in case this is *right*, and choose to go left. This action will be chosen on the same basis as agent 3 made her choice—namely, the deduction of the private signals of agents 1 and 2. The fourth agent will, however, not have a stronger reason to go left than agent 3 had, since the choice made by agent 3 is *uninformative* to all subsequent agents. This is a corollary of agent 3 being in cascade: since agent 4 knows that agent 3 is rational and received the public signal (*left*, *left*), 4 can deduce that 3 would have chosen to go left *no matter what private signal she received*. Hence, agent 4 will base his decision only on the choices of the two first agents, and will also be in cascade. Similar considerations apply to all subsequent agents: they will all be in the cascade, ignoring both their private information and the choices made by previous agents in the cascade.

3.1.3. Theme: Fragility

Given the example run above, it may be seen that *an informational cascade may have a very weak basis*, consisting of only the first two actions in the sequence.[1] This is the reason cascades are often considered *fragile*: the balance in even a long-running cascade may be upset if actions contrary to the herd behavior are observed. If one allows for agents perfectly informed by their private signals in the model described, it will take *only one agent* to break the cascade.

To see this, assume that the fifth agent in the described left cascade *knows* that she should go right instead of following the herd. She would ignore not her own knowledge but rather the public signals sent by previous agents, and therefore choose to go right. Any subsequent agent may now take agent 5's action to indicate that 5 had hard information and then simply choose to follow her instead of the cascade. However, even if agent 5's action is only interpreted as 5 having received a private signal equal in

[1] In case no cascade arises in the beginning of the sequence, one will occur in case there are *two more* agents that choose one action than there are agents choosing the other. Given that agents assume others' signals as likely to be as correct as their own, any cascade will commence on an equally weak basis.

likely correctness to all other agents' private signals, the action is still enough to break the cascade. For now agent 6 will know that agents 1 and 2 received *left* private signals, that agents 3 and 4 were in a cascade, and that agent 5 received a *right* private signal. Hence, 6 will no longer ignore his private signal. In case this is *left*, he will go left, but if it is *right*, he will choose to act in accordance with the tie-breaking rule. Thereby, agent 6 is no longer in the cascade.

3.1.4. Theme: Epistemic Assumptions Regarding the Tie-Breaking Rule

In the example above it was mentioned that agent 2 could deduce the *left* signal of agent 1 and, given that 2 received a *right* signal, he should invoke a tie-breaking rule in order to decide what choice to make. Examples of such tie-breaking rules may be to randomize the available option (Bicchieri and Fukui 1999; Bikchandani et al. 1992), choose an externally given option (Banerjee 1992; Bikchandani, Hirshleifer, and Welch 1992), act in accordance with a private signal/preference (Anderson and Holt 1997; Banerjee 1992; Bicchieri and Fukui 1999), or something else.

The epistemic assumptions regarding the applied tie-breaking rule are seldom made clear. A noteworthy example is the use made by Bikchandani and colleagues when discussing their binary model (in structure identical to the above) and a tie-breaking rule, which requires randomization when indifferent. In the case where the third agent receives the equivalent of a public (*left*, *left*) signal and a *right* private signal, it is noted (as above) that the third agent will be compelled to ignore her private information and go left.

However, if it is generally known that all agents randomize to break ties, then the action signal string (*left, left*) will *not* inform agent 3 of the private signal of agent 2. This follows as agent 2's *left* action might be the outcome of the tie-breaking rule applied to 2's private *right* signal and his deduction of agent 1's private *left* signal. Hence, agent 3's private signal will determine her action, and she will therefore not be in a cascade. Given a *left* signal, she will go left, and given a *right* signal, she will again randomize. This randomization may again result in her going left, putting the fourth agent in his own shoes: though having received a pure string of public *left* signals, agent 4 cannot extrapolate any information other than the signal of the first agent, and will therefore not be in a cascade.

We may consider whether rational agents should assume of other rational agents that they would act in the same manner when indifferent as they themselves would; however, if this is assumed to be the default behavior, then cascades involving a randomizing tie-breaking rule will not require as weak a basis as the above section leads one to believe.

Adopting the assumption that ties are broken by acting in accordance with one's private signal changes this situation. If all agents know that this is the tie-breaking rule used, then this will facilitate later agents' ability to

deduce the correct signal of indifferent agents. One reason for adopting this tie-breaking rule in cascade models is that it fits laboratory evidence far better than randomization (Anderson and Holt 1997 and 2008). As Anderson and Holt note, this assumption is also reasonable in a case where there is a chance that previous subjects have made mistakes in their decisions.

3.1.5. Theme: Positive and Negative Cascades, and Infostorms

Informational cascades are neither good nor bad in themselves—they are merely there. A cascade is the result of rational agents basing their decisions on information extrapolated from the actions of those choosing before them to a degree where this supplies a stronger reason to act than their private information does. Given the assumption that correct private signals are more prevalent than incorrect ones, cascades will in the theoretical framework more often than not herd agents toward the correct choice. In the terminology of Kuran and Sunstein (1999), such cascades are *positive*.

Still, given that there is a risk of the private signal being wrong, there will be a nonzero probability that a *negative* cascade will occur. That is, given that the agents' private signals do not *perfectly* inform them, there is a risk that even rational agents will herd toward the suboptimal choice.

In relation to social media, the weak required basis together with the possibility of negative cascades form a problem for, for example, comment architecture of opinion blogs, product reviews, and so on, as it may be highly valuable for interested parties to hijack the initial segment of a comment thread using sockpuppets[2] in order to form public opinion. Such a use clearly marks how manipulation of public opinion comes into play by the introduction of excessive information. Where multiple sockpuppets are used to voice seemingly similar views, we may fear that this can lead to uncontrolled opinion formations in favor of viewpoints not otherwise supported by the online community. Where rational agents may have good reasons based on higher-order reasoning given the private signals received by others, the use of biased sockpuppets may form opinion based on nothing more than the preferences of one involved party.

The same may be said about the introduction of new products into markets (Easley and Kleinberg 2010, 505). Here, as exemplified by the case of the best-seller list we mentioned in the introduction, later customers observe the choices of previous costumers, but not necessarily their satisfaction with the product. The same goes for real-estate bubbles. Prizes rocket to the heavens because misleading information is spread about the ever-growing value of real estate through the constant bombardment by real-estate brokers, government officials, bankers, and financial advisers. The information ends at Mr. and Mrs. Regular Real Estate Owner, who,

[2] Sockpuppets are fake and misleading online identities used to support, defend, or praise a certain party. The sockpuppet is puppeteered by the party in question but is posed as unaffiliated.

based on the received overwhelming amount of information, end up buying their new house at an unrealististically inflated price. Shiller (2008) calls this *boom thinking*. This amounts to manipulation of agents by providing them with too much misleading information.

Though cascades have been reproduced in laboratory settings (Anderson and Holt 1997), there is reason to doubt that these cascades occurred because the subjects applied Bayes' rule (Huck and Oechssler 2000). Huck and Oeschssler point out that applying Bayes' rule is mentally taxing, most notably in cases where Bayesian updating suggests something different from the counting rule, which suggests simply going with the majority. Huck and Oeschssler note further that the decisions made by subjects in both their own experiments and in the experiments of Anderson and Holt are more often in conformity with "follow your own signal" than with Bayesian updating. This can on the one hand be taken to be good news, since I would seem to suggest that negative cascades are less likely to form in the real world. It may, however, also be taken to suggest bad news: if not even economics students in an exam situation are able to analyze the higher-order reasoning required in a cascade setting sufficiently well so as to not merely act on their own signal, the chances that average Amazon customers will be able to do so is highly unlikely. If it is granted that book shoppers who orientate themselves toward best-seller lists have no or a very weakly believed private signal, we may hypothesize that the opacity of the informational situation at hand will lead them to choose by a simple heuristic, namely, by simply following the salient indicator that many others have previously bought a certain book.

3.2. Bystander Effects

The bystander effect is a notion from social psychology taken to cover the seemingly paradoxical inaction of witnesses in emergency situations where multiple witnesses are present. A paradigmatic example is the story of the murder of Kitty Genovese as referred to by among others Cialdini (2000) and Bicchieri (2006). Bystander effects have frequently been reproduced in laboratory settings (see Latané and Nida 1981 for a review), most notably by Darley and Latané (1968), and multiple explanations have been suggested. Among these are that bystanders believe that others are more qualified to aid than they themselves are, that bystanders feel averse to acting alone in comparison to acting in accordance with a majority, and that bystanders are in a situation of *pluralistic ignorance* resulting in a wrong belief that no help is needed. A notable study utilizing the latter two is Bicchieri and Fukui 1999. Bicchieri and Fukui construct a model involving rational agents in a setting of pluralistic ignorance and (a nonsequential version of) informational cascades to explain the introduction of various unpopular norms, including the bystander effect, college binge drinking, and violent gang behavior.

At the outset in the bystander effect one may be dealing with belief dynamics generating situations involving pluralistic ignorance. The initial focus is on the epistemic assumptions underlying the higher-order reasoning of the social comparison mechanism in play.

3.2.1. Structure

The structure generating this sort of bystander effect includes:

1. a set of agents that act concurrently in a number of rounds,
2. three possible actions in each round, and
3. a preference order in the outcome of choices.

To illustrate the setup in the bystander effect, there may be a set of witnesses in an emergency situation, who act simultaneously in a number of rounds. They can choose to help, not to help, or to inquire or survey further to obtain more information. All agents prefer to help if help is required, but not help otherwise; that is, their preference in choice depends on the true state of the world. If an agent chooses to help or not to help, the agent cannot choose in later rounds. It is, however, cost-free to "skip a round" by inquiring further or surveying the situation.

The decision is again to be performed under uncertainty: agents do not know whether the situation in fact calls for assistance. As with informational cascades, there is no strategic interaction in the decision problem, so no agent will have an incentive to mislead subsequent agents by choosing in contrast to the best of their knowledge. Therefore the choices of other agents can again be interpreted as conveying information regarding others' interpretation of the situation.

Given this, as was the case with informational cascades, agents may choose to base their action not only on their private information but also on the information extracted from their peers. The following epistemic assumptions are made pertaining to the information dynamical structure:

1. the underlying structure is known to all agents,
2. each agent makes a rational decision in each round based on the available information, which consists in
   a. a *public signal* about the true state of the world,
   b. a *public signal* consisting of the actions *performed* by the previous agents,
3. a belief among the agents that others,
   c. given that they believe help is required, are more likely to help, than they are likely to either inquire or not help,
   d. given that they believe help is not required, are more likely to not help than they are likely to either survey or help, and
4. knowledge of rationality as described in 2.

Regarding item 2, three things are to be noted. First, in *a*, agents are assumed to receive a *public* signal about the true state of affairs. This signal consists in the emergency event, for example, a visual impression that an elderly lady falls. This signal is assumed to be *common knowledge*, as everybody can see that everybody else can see the event, and so on. It is not, however, known to other agents how the individual agent *interprets* this signal. Second, agents are not assumed to be made aware by the end of a round whether their actions were in accordance with the true state. That is, no external source of information is available between rounds to inform agents in later rounds. Third, notice the emphasis in *b*: in contrast to the informational cascade case, it is assumed not only that agents perceive the choice, and not the private signal, of other agents, but also that they only perceive the *performed output* of this choice. This is essential, as the choices to survey and to not help are *output equivalent*.

The assumption made in item 3 is that the group of agents already face *pluralistic ignorance* with regard to the decision rules used in the situation. This is a situation "where a majority of group members privately reject a norm, but assume (incorrectly) that most others accept it" (Centola, Willer, and Macy 2005, 1010), but where the norm in question is not a true social norm but rather a decision rule. In conjunction with suitable assumptions regarding payoff and degrees of belief, *every agent will have a propensity to survey the situation instead of helping or not helping*. However, qua item 3, all agents also believe that others reason by a *different* choice rule, namely, that they would choose to help or not help under the same circumstances. To illustrate how this assumption affects agents' interpretation of the public signal, let us consider an example run.

3.2.2. Illustration: Not Initiating a Rescue

To illustrate the bystander effect using a simple setup, consider three agents witnessing an event where an elderly woman trips in the street. Assume that the agents have two rounds in which to decide whether or not to help. The fact of the matter is that the lady needs help. The public signal sent by the event is, however, ambiguous: it may be interpreted as the lady tripping without being hurt or as the lady having badly twisted her ankle. Assume that all agents interpret the signal correctly, and therefore initially believe that the lady requires assistance.

Let us focus on a particular agent, *a*. Given that *a* believes that she is no better at interpreting the public signal than others are, it will be reasonable for her to survey. By surveying, *a* can observe the actions of others, and thereby gather information regarding their interpretation of the public signal. Under the assumption that others are at least as good as herself in estimating the true state from the public signal, this further information will lead to a stronger basis upon which she can later choose to either help or not help.

Notice how the reasoning for choosing to survey implicitly utilizes the assumption of pluralistic ignorance from item 3 above. For *a* to be able to infer information from other agents' actions in the first round, it must be assumed that these actions reflect the agents' private beliefs, even though the action chosen by *a* does not reflect her own beliefs to others.

To see how *a*'s action misrepresents her beliefs to others, recall the assumption in item 2b above, stating that agents perceived the *performed output* of the choices of other agents. In the presented case, the choice to survey and the choice to not help are, however, *output equivalent*: other agents cannot distinguish these two choices from each other, as both outcomes consist in standing still and witnessing the situation at hand. Following the assumption of pluralistic ignorance, all other agents now believe that *a* has chosen *not to help*.

Given that all agents have acted as *a* did in the first round, what new information is *a* left with after she is done surveying the situation? She has seen two other witnesses not doing anything, and as she, due to pluralistic ignorance, believes that they follow a choice rule different from hers, she will infer that they all interpreted the public signal as showing that the true state is one in which no help is required. As this goes for all agents, a situation of *belief-oriented* pluralistic ignorance has occurred: a situation in which "no one believes, but everyone thinks that everyone [else] believes [that no help is required]" (Krech and Crutchfield 1948, 388–89).

As *a* takes the two other witnesses to be her epistemic peers, she will now have compelling reasons for revising her belief. Since the roles of all agents are symmetric, agent *a* is not a special case, though, and hence the second round will commence with all three agents believing that no help is required. As they can obtain nothing from surveying further (as this is the last round), the rational choice will be to *not help*.

In conclusion, a group of rational witnesses suffering under pluralistic ignorance regarding each other's decision rules may by social comparison cause a bystander effect.

3.2.3. Theme: Acting in Conformity

The outlined model for the bystander effect ignores the possibility of agents having *interactive* preferences. If the structure outlined above is conjoined with a preference to act in conformity with a majority, a model for the emergence and persistence of unpopular norms may be constructed (see Bicchieri and Fukui 1999). Though the bystander effect may occur on solely epistemic grounds, as illustrated above, conformity to group behavior plays an important role in situations with a similar structure (see Miller and McFarland 1987).

A good example of how pluralistic ignorance incorporating a preference to conform in a bystander-effect-like setting may have negative consequences is board decisions regarding strategic choices of organizations (see

Halbesleben and Buckley 2004 for a short review of the historical developments of pluralistic ignorance in organizational settings). A round table discussion regarding a strategic choice may easily be seen to have a similar structure: a number of executives will all be witnessing a firm's poor business performance but will fear suggesting that the situation be remedied, due to adverse feelings about acting as a minority and a concern for maintaining the respect of their fellow board members, against a majority who believe that poor performance is due to outside factors, not a current poor strategic choice (see Westphal and Bednar 2005). Based on survey studies, Westphal and Bednar argue that when a firm's performance is relatively low, outside directors will have a tendency to underestimate the degree to which peer directors share their private concerns regarding the status quo strategy, and when less concern is expressed, the tendency toward underestimation increases. They further argue that personal friendship ties seem to diminish this underestimation.

3.2.4. Bystander Effects for Real (Estate)

From around 2003 to the first half of 2007, prices for Danish corporate realty had been going through the roof; especially from 2005 the upward curve for corporate real estate had been extremely steep. General optimism in the market was high, and unrealistically so among the three primary actors in the Danish mortgage market—borrowers, investors (banks, credit institutions, private entrepreneurs), and mortgage merchants. On top of that, unfortunate merry-go-rounds had begun to appear as a result of the easy access to money. Property sharks would take a mortgage deed to the bank to borrow money on it, and often enough the investment would be set up such that the shark could walk away with a loan very much exceeding the value of the deed and property. The shark would then venture to acquire new property, issue new mortgage deeds, take out a loan on those or sell them to other sharks, who would then try their luck with investors for yet more capital to invest. And so it went on for some time while the prices on corporate realty in Denmark went berserk and nobody really checked whether there was value for money. Banks and credit institutions were so eager to issue loans that credit assessment became sloppy to nonexistent. Right before the bubble burst property sharks were mostly trading among themselves at fictitious prices in a world of financial fiction, creating a phenomenon later referred to as *mortgage deed merry-go-rounds*.

Until a Danish newspaper, *Jyllandsposten*, started publishing a series of articles on the mortgage deed merry-go-rounds in late 2007, nobody interfered, although there were plenty of parties who witnessed the unrealistic prices which much corporate realty went for and in which they had a vested interest. Between the Danish National Bank, the Ministry of Business and Growth, the Danish Financial Services Authority, and miscellaneous banks and credit institutions it should have been clear that

regulatory intervention would have been in order long before the bubble came to bursting. But it did not happen, partially for socio-epistemic reasons. Unsure as to whether to do something or at least propose intervention in the carousels, the various bystander parties began observing each other, either so as to become wiser regarding courses of intervening action or because every party thought that the other parties would issue a suggestion as to what to do. Since every party observed every other party all at the same time, nobody did anything, exactly because nobody else did anything. The mechanics was a bystander effect in the mortgage deed merry-go-rounds. Even if one party thought (as a matter of fact, even if all witnesses to the incident thought) that intervention was in order, they could observe all other parties doing nothing. Thus, the Danish National Bank, the Ministry of Business and Growth, the Danish Financial Services Authority, and so forth started subscribing collectively to a norm of nonintervention, which they may very well have privately rejected.[3] That amounts to a state of pluralistic ignorance (Hendricks 2010; Hendricks and Rasmussen 2012).

3.3. Group Polarization

The example involving C. H. Charles's 1920 *Love Letters of Great Men and Women* illustrates a common feat of modern web technologies, namely, inherent *information selection processes*. On Amazon, shoppers are prompted to buy additional items based on what they are currently viewing; on Facebook, the amount of interaction with friends determines their *edge rank* in relation to you, which in turn determines how frequently they appear in your news feed, and Google by default uses your past 180 days' search history to provide Personalized Search for Everyone.[4]

A further common feat of modern web technologies is *social*. Most webpages offer a built-in button to "like," "share," or "comment" on the displayed item. This provides the opportunity to show interest in, or discuss, the content easily on social sites and in the associated comment thread. This allows friends of yours who share your attitude toward a given issue to like the news item and be notified of comments so as to participate in the discussion and reshare it with their social network.

In relation to social discussion, an interesting phenomenon is *group polarization*. Group polarization refers to a reproducible product of group deliberation where each of the group members following a discussion ends up holding a more extreme position regarding some viewpoint than they did prior to deliberation. The phenomenon can reliably be reproduced in

[3] There is evidence to the effect that some players, including the director of the Danish National Bank, on more than one occasion before the bubble burst warned against the overheated corporate realty market, without doing anything about it.

[4] Even when signed out; cf. http://googleblog.blogspot.dk/2009/12/personalized-search-for-everyone.html (accessed 6 January 2013).

lab settings (see Myers and Lamm 1975 and 1976 and Myers 1982 for reviews of experimental literature), using, among others, a setup like the following.

3.3.1. Structure[5]

Group polarization can occur in situations in which there are:

1. a set of agents,
2. an issue on which agents' degree of agreement can vary on a scale with neutral midpoint and two extreme poles,
3. a division of agents into subgroups, which are homogeneous with respect to their degree of agreement relative to the midpoint, and
4. a group deliberation process in which agents are free to discuss their opinions and arguments.

Given one such situation, a subgroup is said to *polarize* or *shift* in case the product of the group discussion has shifted further toward the pole initially favored. The shift is measured by comparing the average degree of individual prediscussion expressions of agreement with a postdiscussion expression. The latter may be given either by asking for postdiscussion expressions from individual agents and finding the mean, or by requesting the group to reach consensus, or by requiring that the group determine this value by majority vote.

Based on homogeneous group experiments much akin to the above in setup, several studies have documented group polarization. Myers 1982 provides an overview of some of these studies. Two examples include racial attitudes among high-school seniors and responses to fictive international military crises involving the United States among U.S. Army officers, ROTC[6] cadets, and university students. In the example regarding racial attitudes among high-school seniors, students were divided into high-, medium-, and low-prejudice groups, and following discussion it was seen that the high and low groups had polarized. The high group had moved from ~1.7 to ~3 on a scale from –4 to 4, with zero being neutral, –4 being low prejudice, and 4 being high prejudice. The low group moved from ~2.8 to ~3.5. In the latter study, groups consisting of, respectively, U.S. Army officers, ROTC cadets, and university students were asked to choose among ten responses ranging from bilateral negotiations to nuclear force. Here, students initially favored the softer responses, whereas officers recommended the more militant solutions. After discussion, these

[5] The following structure is based on one of several experimental approaches described in Myers 1982.

[6] The ROTC is the U.S. Reserve Officers' Training Corps, a college-based military training program.

two groups polarized, whereas the ROTC cadets where more neutral in both prediscussion and postdiscussion scores.

3.3.2. The Black Box of Group Discussion

The main task in explaining the general phenomenon of group polarization consists in unpacking the black box of group deliberation leading to an opinion shift (see Myers 1982; Isenberg 1986). One suggested explanation focuses on *informational influence*. According to this theory, subjects in the deliberation processes receive and weigh information that affects their opinion on the issue at hand. It is assumed that the initial lean in direction influences the number of arguments pro and con the given direction in favor of the leaned-to pole, and that more arguments in favor of the initial lean are therefore presented. Given that not all arguments have been considered by all agents, some agents will become more convinced of the leaned-to direction, thereby shifting the mean opinion of the group toward the given pole.

Several studies indicate that there is a certain structure to the arguments that provide a shift in opinion. Myers and Bishop 1971, Bishop and Myers 1974, and Vinokur and Burnstein 1974 suggest and support the view that the group shift is based on a number of parameters, namely, the direction of argument (which pole the group favors), the cogency or perceived validity of the argument, and the argument's novelty (the degree to which the argument was new to agents in the discussion).

By way of example, assume a homogeneous group of three agents initially agreeing on some stance to degree 2 on a scale from –4 to 4 because they each recall two arguments in favor of the positive direction. During discussion, they all advance their arguments, each hearing one novel argument from either of the other agents, one of which they find convincing. Assuming that each argument affects their degree of agreement by 1, each agent will after the discussion have changed their degree of agreement to 3, thereby producing a group attitude shift of 1.

The informational influence approach explains why already quite polarized groups show less polarization than do initially less polarized groups—the individuals in the former are in the initial state closer to the distributed "knowledge" fix point.

It has been argued (Myers 1982) that an additional element of *argument rehearsal* in group discussions amplifies the belief formation in groups, thereby creating a stronger polarization effect. This is supported by findings to the effect that being passively presented with arguments in favor of a direction does not produce as large a shift as active discussion does. Instead, arguments need to be rehearsed and internalized in order for an attitude change to have proper effect. Myers 1982 further proposes that discussion prompts agents to take a more one-sided line of argument, whereas solemn contemplation elicits a more diverse approach.

3.3.3. Social Comparison

In addition to the information influence theory, it has been suggested that group polarization occurs as a product of *interpersonal comparison*. The main idea behind this approach is that people in groups seek to represent themselves in a favorable light. In order to obtain this goal, subjects must obtain a notion of what qualities are desired in the present group, and then act in accordance with these qualities. If subjects further have a wish to present themselves not just in a favorable light but in a more favorable light than others in the very same group, and all act according to this maxim, the group will elicit a shift in action toward what is viewed in the more favorable light (Myers and Lamm 1975 and 1976; Myers 1982; Isenberg 1986).

There are at least two variants of social comparison theory (Isenberg 1986), one favoring the dissolution of pluralistic ignorance through group discussion, and one favoring bandwagoning effects.

The pluralistic ignorance explanation suggests that the viewpoints presented by agents strike a compromise between the agents' personal ideals and the perceived norm of the group. Given that all agents have initially underrated the degree of agreement with the given direction, group discussion may gradually dissolve this pluralistic ignorance, and a general opinion shift can be observed. This viewpoint is referred to as "release theory" in Myers 1982, as the discussion "releases" agents from the hold of pluralistic ignorance.

The driving force in bandwagoning explanations is that agents wish to "one-up" others. Given that one pole is thought to be the favored viewpoint of the group, an agent may claim to hold a position slightly more extreme than the perceived mean. Through discussion, agents' perceptions of others' positions move in the favored direction, which in turn means that the individual will adopt a slightly more extreme position in the final evaluation.

In a meta-analysis of collected data, Isenberg 1986 concludes that the information influence and social comparison approaches describe conceptually independent processes, which in most cases co-occur. The processes have, however, been shown to exist in isolation. Experiments in which social comparison cannot be used because agents were either unaware of each other's initial choices or were not informed of the eventual scale on which the group decision were to be assessed still demonstrate polarization effects, which lend credit to the information influence explanation. However, situations in which only the attitudes of others are made available and discussion is eliminated also show group shift, lending credit to the social comparison theory (see Myers 1982).

3.3.4. Online Information Selection and Homophily

We are daily surrounded by information selection processes that provide us with information akin to what we have previously searched for or liked.

With respect to group polarization and information influence, we fear that living in such *filter bubbles* (Pariser 2011) provides logs to fuel the fire of group polarization, since such information selection processes provide us primarily with arguments in favor of our initial views. Pairing this up with the further information selection processes of both natural and artificially enforced homophily, produced by mechanisms such as edge ranking on social network sites, may further produce stronger tendencies to share, like, and discuss events with like-minded others, allowing both informational influence and social comparison dynamics to run out of control.

4. "Weather Conditions" for Infostorms

The previous section described three different types of social information phenomena: informational cascades, bystander effects, and group polarization. These may affect individual beliefs and actions in large groups of agents in ways that track truth imperfectly, so that people end up believing false propositions and as a result possibly act contrary to their goals and interests when faced with uncertainty. These phenomena have been studied empirically for years in social psychology, but formal approaches such as logic, game theory, decision theory, and the general analytical approaches well known to philosophers offer rigid frameworks for analyzing the exact components involved in derailing otherwise rational agents from the truth. From this perspective, social information phenomena become just as interesting from a philosophical point of view as more classic areas of epistemology that focus on how individual agents may obtain justified true beliefs in nonsocial settings.

Rigid philosophical approaches to social information phenomena may turn out to be particularly fruitful to society if applied to the many information technologies and systems devised to deal with information overload in "the age of abundant information." This could be of great benefit to democratic society at large, not only because it seems clear that these technologies and systems play a crucial role in boosting problematic social information phenomena so that they may be characterized as infostorms, but also because it seems that these technologies and systems could be designed according to insights obtained from rigid epistemological approaches to social information phenomena with their emphasis on truth tracking and truth preservation.

As an example of a simple and well-known information system, think of the ever-expanding market of books (partially caused by the decreased cost and complexity of producing this good), where best-seller lists have seized the role of a central information technology for guiding individual decision-making under uncertainty analogous in function to the heuristics often used in individual decision-making. Used by the consumer, the best-seller list exhibits several of the information phenomena discussed above. For one, it is used as a device to solve a general decision problem

generated by a situation where agents have too little information and thus look to other evidence to inform their impeding decision. However, the nature of the best-seller list is to compile overall sales in sampled bookstores from which a list is created showing the best-selling books ranked in order of sales. Accordingly, a best-seller list embodies a social information process, where consumers faced with the uncertain decision as to which books are worth reading turn to this list for information, but usually forget that sales figures reflect aggregate buying behavior, rather than anything about the quality of the book bought.

Viewing online book shopping as a decision problem with *n* actions (books) that the buyer may choose from and assuming the book market to be driven by a series of buyers who each take turns in buying a book, we see that best-seller lists are likely to produce informational cascades. Each agent in the series of buyers chooses to buy a book, where the sale is used as input for the best-seller list facing the next buyer. Hence, if initial movers in this series all buy a particular book, then others are likely to follow, thereby producing a cascade as discussed in section 3. In fact, it was informal knowledge about the structure and dynamics of informational cascades that allowed the two authors of *The Discipline of Market Leaders* to manipulate their way to a high ranking on the *New York Times* best-seller list and a tidy fortune.

It is also interesting, however, to notice that besides exhibiting an informational cascade, a best-seller list exhibits pluralistic ignorance, since each buyer looks to social evidence in a situation of uncertainty, failing to see that the evidence he is witnessing might very well be the outcome of a search process identical to his own, rather than an expression of qualitative evidence. Thus, best-seller lists exhibit the interesting phenomenon of reproduction of pluralistic ignorance by an informational cascade as well as exemplifying the wider point that the information phenomena discussed in this article often interact.

While the best-seller list is not a paradigm example of an infostorm in itself, it is obvious that when systems similar to the best-seller list are implemented for directing the choices of large masses of agents on the Internet, the number of people affected becomes ever larger and the consequences ever more severe. *Love Letters of Great Men* serves as one example of this, but to the extent that search engine optimization (SEO) ranking is determined by the number of clicks a site gets, the resulting infostorm is identical in informational structure to that of the best-seller list. People will click on sites when they feature high in the SEO ranking, because they believe that this indicates that these sites are valuable, albeit the ranking only indicates that other people have clicked on these pages as well. Furthermore, this pluralistic ignorance may be reproduced on a higher level as the belief that other people actually find interesting what is found on these pages and hence that there is something interesting on them.

Ultimately, this phenomenon may be seen as a crucial characteristic of the holy grail of modern online marketing: "virality." If you can get enough people to click on a particular YouTube video or link, the video or link will be spread through huge numbers of agents—it will go "viral." While the fact that these links are usually consulted and the videos are watched before they are spread further adds a bit of complexity, it doesn't detract from the phenomenon just described. That people believe that other people actually find what is found on these pages interesting may ultimately override the private signal that the link or video is actually dull or stupefying, and lead the agent to share this link or video anyway in the hope of spreading an interesting link or video to her social network.

In the age of information overload, however, clicking on links may even engage with well-intended algorithms that end up inadvertently producing infostorms in the attempt to handle this overload by filtering out irrelevant information. Thus, for instance, Facebook's Edge-rank algorithm introduced in 2009 and Google's individualized search feature as well as programs such as Zite may easily create so-called filter bubbles without users ever even noticing (Pariser 2011). Basically, these filter bubbles have close shaves with group polarization—where the group in fact consists of a single real agent—through being constructed by algorithms which are trying to serve only what is relevant to you, but which ultimately amount to serving what fits your worldview as revealed by yourself through your clicking behavior. While your clicks will show a confirmation bias toward your leanings, the online world will come to reflect this bias. In turn, this phenomenon of belief polarization or singleton-group polarization may also hook up with pluralistic ignorance. You will be looking to Facebook, Google, or Zite as honest conversation partners providing you with information about the world, while these algorithms are actually trying to select and feed you with information that you like.

5. Information Control Problems

How to deal with infostorms is an open problem. So far nothing has been said about how rigid analysis may help to undermine their negative effects. Yet, information-driven social epistemic phenomena may be viewed like control problems in engineering. Initially the structure of the problem should be revealed and understood. In fact, even the informal structural characterization of informational cascades, pluralistic ignorance, and bystander effects provided above may be adapted to a formal logical setting exactly put in play to account for the informational processes and epistemic dynamics between interacting agents. Private and public signals, preferences, actions, events, scales, and weights are all crucial components isolated so far to account for the structure and epistemic mechanics of the information phenomena in question.

By way of example, Rendsvig and Hendricks (forthcoming) scrutinize investment behavior in the light of the skip, gamble, or quit game, in relation to the following narrative. Any investor, especially in the wake of the current situation of the financial markets, is faced with a difficult investment problem: Should I skip, gamble, or quit? Uncertain as to whether to skip, gamble, or quit, in order to become wiser the investor starts looking around at other investors to see what they are doing. Other investors may be looking back because they are also unsure what to do, as they are likewise short of decisive information. Investors may start looking for *social proof* to facilitate a qualified decision. Given social proof, for the individual investor skipping, gambling, or quitting all of a sudden becomes contingent upon information about what the investor expects about a market crash, what other investors are expected to do based on their expectations about a market crash, and whether the other investors are (believed to be) aggressive or conservative with respect to their financial behavior. It also means that the collective behavior of investors become susceptible to the workings of socio-epistemic phenomena like informational cascades, pluralistic ignorance, and bystander effects.

This makes investment behavior essentially an informational control problem of social proof, including:

- entities like agents, actions, expectations, modes of behavior,
- modeling tools like epistemic logic, game theory, and judgment aggregation, and
- parameters like uncertainty, available information, decision rules as to what to do personally, interpretation rules of other agents' behavior, and social network structure.

Control problems in engineering and technology are often *modular* in the sense that twitching or shifting entities, tools, or parameters changes one control problem into another control challenge. The same is the case here. It turns out that one may go from the study of informational cascades to the study of bystander effects by changing some modules while retaining others, plugging the modules, and pressing "play" (Hendricks and Rendsvig in progress). And so on for other socio-epistemic phenomena like bandwagon effects, boom thinking, conformity, compliance, gullibility, opinion bubbles. . . . We have only just begun to play.

Pelle G. Hansen
The Initiative for Science, Society & Policy
University of Southern Denmark and Roskilde University
Universitetsvej 1
DK4000 Roskilde
Denmark
pgh@ruc.dk

Vincent F. Hendricks
Department of Philosophy
University of Copenhagen
Njalsgade 80
DK2300 Copenhagen
Denmark
vincent@hum.ku.dk

Rasmus K. Rendsvig
Department of Philosophy
University of Copenhagen
Njalsgade 80
DK2300 Copenhagen
Denmark
rendsvig@gmail.com

References

Anderson, L. R., and C. A. Holt. 1997. "Information Cascades in the Laboratory." *American Economic Review* 87, no. 5:847–62.

———. 2008. "Information Cascade Experiments." In C. R. Plott and V. L. Smith (eds.), *Handbook of Experimental Economics Results*, volume 1, 335–43. North Holland: Elsevier.

Banerjee, A. V. 1992. "A Simple Model of Herd Behavior." *Quarterly Journal of Economics* 107, no. 3:797–817.

Bicchieri, C. 2006. *The Grammar of Society: The Nature and Dynamics of Social Norms*. New York: Cambridge University Press.

Bicchieri, C., and Y. Fukui. 1999. "The Great Illusion: Ignorance, Informational Cascades, and the Persistence of Unpopular Norms." *Business Ethics Quarterly* 9, no. 1:127–55.

Bikhchandani, S., D. Hirshleifer, and I. Welch. 1992. "A Theory of Fads, Fashion, Custom, and Cultural Change as Informational Cascades." *Journal of Political Economy* 100, no. 5:992–1026.

Bishop, G. D., and D. G. Myers. 1974. "Informational Influences in Group Discussion." *Organizational Behavior and Human Performance* 12:92–104.

Centola, D., R. Willer, and M. Macy. 2005. "The Emperor's Dilemma: A Computational Model of Self-Enforcing Norms." *American Journal of Sociology* 110:1009–40

Cialdini, R. C. 2000. *Influence: Science and Practice*. New York: Allyn and Bacon.

Cooper, J., K. A. Kimberly, and K. Weaver. 2004. "Attitudes, Norms and Social Groups." In M. B. Brewer and M. Hewstone (eds.), *Social Cognition*, 234–56. Oxford: Blackwell.

Darley, J. M., and B. Latané. 1968. "Bystander Intervention in Emergencies: Diffusion of Responsibility." *Journal of Personality and Social Psychology* 8:377–83.

Easley, D., and J. Kleinberg. 2010. *Networks, Crowds, and Markets: Reasoning About a Highly Connected World*. Cambridge: Cambridge University Press.

Halbesleben, J. R. B., and M. R. Buckley. 2004. "Pluralistic Ignorance: Historical Development and Organizational Applications." *Management Decision* 42, no. 1:126–38.

Hansen, P. G., and V. F. Hendricks. 2011. *Oplysningens blinde vinkler*. Copenhagen: Samfundslitteratur. Forthcoming in English as *Infostorms: How to Take Information Punches and Save Democracy*. New York: Copernicus Books.

Hendricks, V. F. 2010. "Knowledge Transmissibility and Pluralistic Ignorance." *Metaphilosophy* 41, no. 3:279–91.

Hendricks, V. F., and J. L. Rasmussen. 2012. *Nedtur! Finanskrisen forstået filosofisk*. Copenhagen: Gyldendal Business.

Hendricks, V. F., and R. K. Rendsvig. In progress. "Structures of Social Proof."

Huck, S., and J. Oechssler. 2000. "Informational Cascades in the Laboratory: Do They Occur for the Right Reasons?" *Journal of Economic Psychology* 21, no. 6 (December): 661–71.

Isenberg, D. J. 1986. "Group Polarization: A Critical Review and Meta-analysis." *Journal of Personality and Social Psychology* 50, no. 6:1141–51.

Kahneman, D. 2011. *Thinking, Fast and Slow*. New York: Farrar, Straus and Giroux.

Katz, D., and F. H. Allport. 1931. *Student Attitudes*. Syracuse, N.Y.: Craftsman.

Krech, D., and R. S. Crutchfield. 1948. *Theories and Problems of Social Psychology*. New York: McGraw-Hill.

Kuran, T., and C. Sunstein. 1999. "Availability Cascades and Risk Regulation." *Stanford Law Review* 51:683–768.

Latané, B., and S. Nida. 1981. "Ten Years of Research on Group Size and Helping." *Psychological Bulletin* 89, no. 2:308–24.

Miller, D. T., and C. McFarland. 1987. "Pluralisitic Ignorance: When Similarity Is Interpreted as Dissimilarity." *Journal of Personality and Social Psychology* 53:248–305.

Myers, D. G. 1982. "Polarizing Effect of Social Interaction." In H. Brandstätter, J. H. Davis, and G. Stocker-Kreichgauger (eds.), *Group Decision Making*, 127–60. New York: Academic Press.

Myers, D. G., and G. D. Bishop. 1971. "The Enhancement of Dominant Attitudes in Group Discussion." *Journal of Personality and Social Psychology* 20:386–91.

Myers, D. G., and H. Lamm. 1975. "The Polarizing Effect of Group Discussion." *American Scientist* 63, no. 3:297–303.
———. 1976. "The Group Polarization Phenomenon." *Psychological Bulletin* 83, no. 4:602–27.
Pariser, Eli. 2011. *The Filter Bubble: What the Internet Is Hiding from You*. New York: Penguin.
Rendsvig, R. K., and V. F. Hendricks. Forthcoming. "Social Proof in Extensive Games." *Studia Logica*.
Shiller, R. J. 2008. "How a Bubble Stayed Under the Radar." *New York Times*, Economic View: Business Review (January 15).
Smith, L., and P. N. Sørensen. 2011. "Observational Learning." In S. N. Durlauf and L. E. Blume (eds.), *The New Palgrave Dictionary of Economics*. Online edition. Basingstoke: Palgrave Macmillan.
Sunstein, C. R. 2009. *Going to Extremes: How Like Minds Unite and Divide*. Oxford: Oxford University Press.
Vinokur, A., and E. Burnstein. 1974. "Effects of Partially-Shared Persuasive Arguments on Group-Induced Shifts: A Group-Problem-Solving Approach." *Journal of Personality and Social Psychology* 29:305–15.
Westphal, J. D., and M. K. Bednar. 2005. "Pluralistic Ignorance in Corporate Boards and Firms' Strategic Persistence in Response to Low Firm Performance." *Administrative Science Quarterly* 50:262–98.