For *Routledge Handbook of Philosophy of Information* / March 3 2015

The Philosophy of Distributed Information

Vincent F. Hendricks

CIBS - Center for Information and Bubble Studies, University of Copenhagen, Denmark

Rasmus K. Rendsvig

Dept. of Theoretical Philosophy, Lund University, Sweden / CIBS - Center for Information and Bubble Studies, University of Copenhagen, Denmark

1. Information in Groups and Social Proof

As an individual you possess a lot of information or knowledge, but groups may be informed or knowledgeable as well. Sometimes information or knowledge in groups is more epistemically potent and conducive to deliberation, decision and action than individual possession – and then again sometimes not.

On June 5, 1989 a young student, later nicknamed the "Tank Man" or "Unknown protester", armed only with two plastic bags was able to stop a column of Chinese tanks during the Tiananmen Square protests. How is that possible? During the Tiananmen Square protests, the Chinese regime was well aware of the threat which thousands of unarmed students and workers could present against a comprehensive political and military machine by merely sitting down in the public space. Thus, the regime sent soldiers and tanks into the square. Soon they had to realize that even a single man with two plastic bags in his hands can stop a column of tanks, as one cannot mow down an unarmed protester if *only everybody watches, and everyone knows that everyone is watching, everyone knows that everyone knows that everyone is watching and so on*. This sort of knowledge is formally referred to as *common knowledge* and may, as the example illustrates be epistemically very commanding, but also quite important for explaining a wide range of other agent interactions, from bilateral trade to notions of communality and rational opinion aggregation.

Another type of group knowledge, distinct from common knowledge, is referred to as distributed knowledge. Knowledge may be unevenly distributed over the members in a group – Bob knows one thing, while Alice knows another, and together they know more, so distributed knowledge may roughly be characterized as the sum of all the knowledge that a collective has available for solving a problem, reaching a decision or performing an action to some desired end. Now, both distributed and common knowledge notions are tracking the truth for groups of agents. And while this tracking feature, individually or collectively, indeed characterizes knowledge, agents with mere information may ever so often get thrown off the truth track.

Agents acquire information from at least two sources: from their immediate environment and what their senses dictate, and from what other agents apparently decide to believe or do. In case of uncertainty as to what to believe or do, individual agents try to tap the immediate environment for more information to become wiser or facilitate qualified decision. But when the environment has no more information to offer or, for some reason, bars additional tapping, agents may decide to consult or observe other agents. This latter source of information is known as *social proof* in social psychology (Cialdini, 2007) and may be an extremely influential vehicle for deliberation, decision and action individually and jointly (Hansen, Hendricks, Rendsvig, 2013): *Single agents assume beliefs, norms or actions of other agents in an attempt to reflect the correct view, stance, behavior for a given situation*.

Sometimes social proof gives the right guidance, other times wrong guidance as to what to think or do among agents in a group. Social psychology and information theory have documented a number of socio-informational phenomena relying on social proof in which agents get more confused and off the truth track than rationally aligned by following the beliefs, norms of other agents (Hendricks and Hansen, 2014). The socio-informational phenomena range from bystander effects, cascades and bandwagons, to belief polarization, all of which, together with common and distributed knowledge, may be formally characterized and their dynamics accounted for.

2. Informational Attitudes in Groups

Common knowledge is a *group notion*, in the sense that it pertains to the knowledge held not by individual agents, but a group of agents. Common knowledge is the strongest notion of group knowledge in the literature. Several weaker notions are useful and important. To describe more precisely these notions of group knowledge, what is needed is:

1. A set of agents capable of possessing knowledge about some set of ground facts
2. Assignments of knowledge to groups of agents
3. The notion of *higher-order information*

To denote agents let us use *a,b,c,...*, and let $K_a A$ be read "*a* knows that *A*", where *A* is a proposition that is either true or false (see Ch. 8), like "Alice's thesis defense is in Room 2-02".

2.1 Distributed Knowledge

Let $G = \{a,b,c\}$ be a group of agents. What would make "the group *G* knows that *A*" true? One candidate could be that *somebody* in *G* must know *A*: either $K_a A$ or $K_b A$ or $K_c A$ must be true. At least if this is the case, then the group may make it to Alice's defense, as Alice may inform Bob and Carol about where to go. This is an instance of *distributed knowledge*: the knowledge that *A* is distributed across the members of the group, so that when the group pools its informational resources, it knows the proposition. Formally, *G* has distributed knowledge of *A* is written $D_G A$.

That Alice must *privately* know that *A* is not required for distributed knowledge. It could, for example, be the case that Bob knows that the defense is either in Room 2-02 or in Room 1-02, while Carol knows that it is in either 2-02 or 3-02, while Alice, being on the brink of a nervous collapse, has forgotten to check. In this case, it is still distributed knowledge in *G* that Alice's defense is in Room 2-02: if Bob and Carol were to pool their information, only one possibility would be left. Hence, a group may have distributed knowledge of propositions which none of its members privately know. Moreover, adding members to a group only increases the distributed knowledge: though *A* is in fact distributed knowledge already in $\{b,c\}$, adding Alice to obtain $G = \{a,b,c\}$ does not mean that the distributed knowledge is lost.

If *G* has distributed knowledge that *A*, does it then mean that they know where to go? Not necessarily. A group has distributed knowledge of *A* in the case that *if the group members were to share all their information with each other, then they would individually know A*. It is not part of the definition that Alice, Bob and Carol in fact communicate. In this sense, distributed knowledge is *potential knowledge* of the group.

The most celebrated instance of distributed knowledge is the *wisdom of crowds* in which the aggregated knowledge of a group is epistemically superior to the performance of the singular agents making up that group (Surowiecki 2004).

2.2 Everybody Knows That

If "group *G* knows that *A*" is to have a stronger meaning than that *A* is "merely" distributed knowledge among *G*'s members, then a natural candidate for a definition is that *everybody* in *G* knows *A*. This knowledge type is often formalized using the "Everybody in G knows that"-operator

E_G , where the proposition $E_G A$ ("everybody in G knows that A") is true just in case $K_a A$ is true and $K_b A$ is true and so forth, for all members of G.

Clearly, if $E_G A$ is true, then $D_G A$ is true, i.e. if everybody in G knows that A, then A is also distributed knowledge in G. Hence "everybody knows that"-knowledge is stronger than distributed knowledge, but is it strong enough for Alice, Bob and Carol to make it to the right room? Each has the information to make it there, as it follows that they privately know the room number. Hence, finding the room is not contingent on information sharing and Alice is safe. But what if Bob and Carol will only go if they know that the other will go as well? Is "everybody knows that"-knowledge strong enough for Bob and Carol to *knowingly coordinate* on going to the right room? Here the answer is *No*. Though both know where the room is, both may be in doubt about *whether the other knows where the room is*, and hence about whether the other will show up. It is perfectly consistent with "everybody knows that"-knowledge of A that Bob lacks information about *what information Carol possess*, for which reason he will not go. Hence everybody knowing is not necessarily sufficient for a coordinated effort of getting to the room in concert. Similarly, a broken car will not move an inch unless the participants coordinate the effort of pushing it all together at the same time. Not only do they have to know this fact each one of them, they also have to know that the others know that they know … . But that is higher-order information and quite different from everybody knowing whatever it is they all individually know.

2.3 Higher-Order Information

In one sense, both distributed knowledge and "everybody knows that"-knowledge are simple notions. Both "only" involve describing the knowledge agents and groups of agents have about ground facts, that is, facts that do not involve knowledge, belief or other propositional attitudes (desire, intention,…). Such propositions, like A about the room number of Alice's defense, are said to contain only *zero-order information*. Building from this, the proposition $K_a A$ is said to be *first-order information*: in referring to Alice's *knowledge*, it contains information about *one level* of propositional attitudes towards a ground fact. Similarly, the propositions $K_a K_a A$ and $K_b K_a A$ are *second-order information*, as they contain two levels of propositional attitudes, $K_b K_a K_a A$ is *third-order*, etc. In short, *higher-order information* refers to second-order information and above. Characteristic of higher-order informational propositions is thus that they contain information about an agent's information about some (other or same) agent's information.

2.4 Common Knowledge

If Bob will only show up if he knows that Carol will show up, and Carol will only show up if she knows that Bob will show up, and both of them know this about each other, then how many levels of higher-order information is needed before they show up to support Alice? The answer may be surprising: infinitely many. Let A be the proposition "Alice's thesis defense is in Room 2-02", and assume that K_bA. This is not enough for Bob to show up, as he will only go if Carol also goes. Assume that also K_bK_cA. Is this enough? Bob now knows that Carol can find her way, but does he know that she will show up? No, for the same reason that K_bA was not enough for Bob. But what if $K_bK_cK_bA$? Then Carol knows that Bob knows where the exam is, and Bob knows this! Surely they must coordinate! Under ordinary circumstances, most likely they would. But if one is strict about the setup, then there is still room for error. In particular, as K_bK_cA was not enough for Bob to go, then K_cK_bA is not enough for Carol either. As Bob tolerates no room for error, knowing K_cK_bA is therefore not enough. We may keep adding additional levels of knowledge, but it will not be sufficient: as long as we add only finitely many, then Bob will always consider it possible that Carol does not know enough about his intentions to go, and will hence not get on the bus. If this seems excessive, consider that the same logic applies to contexts in which there must be zero-tolerance for any coordination failure, such as potential nuclear conflicts.

The problem of coordination will be solved if Bob and Carol share the strongest form of group knowledge, namely *common knowledge*. A group G has common knowledge of a proposition A if everybody in G knows that everybody in G knows that … everybody in G knows that A, for *all* higher-order levels. That G has common knowledge that A is written in notation as C_GA, and may be defined using the E_G-operator: let E_G^1A denote E_GA and let $E_G^{k+1}A$ denote $E_GE_G^kA$. Then C_GA if, and only if, E_G^nA for all natural numbers n. As there are infinitely many natural numbers, common knowledge incorporates an infinite hierarchy of higher-order information.

Common knowledge suffices for resolving the coordination problem between Bob and Carol because Bob's doubt about Carol will vanish as the possibility of error has been replaced by infinite assurance on all higher-order levels. And that goes for Carol too, the same way it went for the "Unknown protester" and the driver of the tank on Tiananmen Square. It's common knowledge now.

2.5 Pluralistic Ignorance

Did you ever go to a show because your friends seemingly wanted to? Did it ever turn out that you all would have preferred to stay at home? If so, then you may have been in a state of *pluralistic* *ignorance* when you made your decision. You collectively get to subscribe to a norm that you privately reject in part because you incorrectly believe that everybody else believes something although no one believes it as it were.

Pluralistic ignorance is an evil cousin of common knowledge. Where common knowledge is truth tracking and can make groups act as one, pluralistic ignorance uses the fact that groups are composed of individuals that seldom know each others' exact thoughts to create highly inefficient groups.

As common knowledge, pluralistic ignorance is also a higher-order notion, but defined using belief: A group is in a state of pluralistic ignorance with respect to proposition *A* if

1. All members of the group believe *not A*,
2. All members of the group mistakenly believes that everybody else believes *A*. (Kretch & Cruthfield 1948), (Hansen, Hendricks, Rendsvig 2013), (Bjerring et al. 2013)

Pluralistic ignorance is unfortunately far easier to achieve than common knowledge, and it plays tricks on us regularly. In some cases, the result is harmless: in attempting to accommodate the preferences of the other, dating couples might go see movie *A*, where both would have preferred to see movie *B*. In other cases, the result may cause liver deficiency: if everybody on campus seemingly enjoys heavy alcohol consumption, new students may attempt to fit in by following suit (Prentice and Miller, 1993).

3. Social Proof in Action

Informational attitudes, like distributed knowledge, common knowledge and pluralistic ignorance describe *static* situations of knowledge and belief. Through observation and communication such static situations may change. Such changing situations are described by *information dynamics* (van Benthem, 2014)

3.1 The Bystander Effect

One dynamic situation type in which pluralistic ignorance is the lead character is known as the *bystander effect*. The phenomenon covers 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 by among others Latené and Darley (1970). In a case of much heated debate from New York City in the 1960s, the 28-year-old Catherine “Kitty” Genovese was assaulted and stabbed on the stoop to her front door. It happened despite scores of neighbours who witnessed large parts of this horrific chain of events, which lasted over half an hour. Subsequently the press reported that no less than 38 witnesses had admitted that they had omitted to act or call the police. In the public debate that followed, the common reader had no doubt as to what the explanation was. Like any other metropolis, New York City had made its citizens callous and indifferent towards fellow citizens. Looking closer at the press reconstruction of the neighbours' own explanations, it was however the fact that no one else seemed to have reacted, that had caused people to refrain from acting. The lack of reaction had instead made everyone believe that it wasn't a case of definite assault, but rather two lovers quarrelling. In other words, it was the ambiguity of the situation coupled with confusion over responsibility and pluralistic ignorance, which lead to this tragic example of the bystander effect. That is, when others don't react, the individual views this information as a sign that a reaction is neither required nor socially demanded. The sad point at the end of the day *is that no one does anything, precisely because no one does anything.*

Bystander effects have frequently been reproduced in laboratory settings[1] , 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 state of pluralistic ignorance resulting in a wrong belief that no help is needed. Here, we focus on the latter explanation (see Rendsvig (2014) for informational dynamics models; (Bicchieri & Fukui, 1997) for game theoretic models).

3.2.1. Pluralistic Ignorance in the Bystander Effect: Structure

The structure generating this sort of bystander effect includes:

1. a state of nature that determines whether an emergency has occurred or not,
2. a set of agents that act concurrently in a number of rounds,
3. three possible actions in each round, and
4. a preference order on 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

1 See Latané and Nida (1981) for an overview paper.

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. The decision is performed under uncertainty: agents do not know whether the situation in fact calls for intervention. 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. Hence, if agents are in doubt about what to do, surveying the situation may seem like a good choice, as it will allow the agent freely to gather additional information on which to base their decision.

There is no strategic interaction in the decision problem, so agents have no incentive to mislead others by choosing in contra to the best of their knowledge. Therefore the choices of other agents can be interpreted as conveying information regarding others' interpretation of the situation. Given this, agents may choose to base their action not only on their private information but also on the information extracted from their peers, i.e. on social proof.

The following epistemic assumptions are made pertaining to the information dynamical structure:

1. the structure above is known to all agents,
2. common knowledge that each agent makes a rational decision in each round based on the available information, which consists in
    a. a *public signal* indicating the true state of the world,
    b. a *public signal* consisting of the *actions performed* by the 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, and
    d. given that they believe help is not required, are more likely not to help than they are likely to either survey or help,

Pertaining to item 2, note three things: 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, of 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 each individual agent *interprets* this signal. Second, agents are not assumed to being 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*: it is only assumed that agents perceive the *performed output* of others' choice, not the choice itself. This is essential, as the choices to survey and not to help are *output equivalent*: a person inconspicuously looking around looks very much like someone not helping.

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.

The assumption of pluralistic ignorance is crucial. Though the decision rules of all agents give them a propensity to survey the situation when in doubt, they simultaneously believe that others reason by a *different* choice rule, namely, that they would choose to help or not to help under the same circumstances. To illustrate how this assumption affects agents' interpretation of the public signal, let us go through the dynamics.

3.2.2. Not Initiating a Rescue

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.

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 deducing the true state from the public signal, this further information will lead to a stronger basis upon which she can subsequently choose to either help or not help.

Notice how the reasoning for choosing to survey implicitly utilizes the assumption of pluralistic ignorance. 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 not to help are *output equivalent*: other agents cannot distinguish these two choices from each other, as both outcomes consists 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 decision rule different from her own, 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 new situation of pluralistic ignorance arises: after surveying, all agents believe that an accident occurred and believe that everybody else believes that nothing requiring intervention happened!

As *a* takes the two other witnesses to be her epistemic peers, she will now have compelling reasons for revising her belief: she will change her mind, and conclude that her initial interpretation was wrong, and now believe that no intervention is required. 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 *not to help*.

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

3.2.3. 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, 1997). Though the bystander effect may occur on solely epistemic grounds, as illustrated above, conformity to group behaviour plays an important role in situations with a similar structure (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 in board decisions regarding strategic choices of organizations (Halbesleben and Buckley 2004). 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 (Westphal and Bednar 2005). Bystander effects even occur in situations with big institutional agents, such as banks, credit institutions and private entrepreneurs. An example of this is the price increase of Danish corporate realty between 2003 and 2007 (Hendricks and Rasmussen 2012), (Hansen, Hendricks, Rendsvig, 2014).

3.2 Cascades

Did you ever go to see a movie because several friends had told you that they had heard it was good? Or buy a book because it was high on a best-seller list? Or choose one restaurant over another because it had more customers? If so, then you might have been part of a *cascade*.

A cascade may metaphorically be compared to a domino effect in a population: when all in an initial group make the same choice, others may choose to follow suit, reasoning that the initiators must each have had good reasons for their choice. Hence, the aggregated choices of the initiators send a public signal that their choice is a good one, worth following. Once more follow suit, and this signal only grows stronger.

There may be good reasons to trust such social proof: people do often make informed decisions, and when there is insufficient time to survey the available options, using social proof as an aggregation method may indeed prove fruitful. An example of a rational cascade, an *informational cascade*, stems from the seminal paper of (Bikchandani, Hirshleifer, Welch, 1992). However, as in the case of the bystander effect, social proof may also lead to unfortunate outcomes, as we are warned against in the age-old retort "If all your friends jumped off a bridge, would you do it, too?"

3.2.1 The Structure of Cascades

In general terms, the structure underlying rational cascades consists of

1. a state of nature, determining a fact in relation to which one must act,
2. a set of rational agents that act sequentially,
3. a set of options between which the agents can choose, and
4. a preference order on the outcome of each choice, in relation to the state of nature.

Let us use the example of a restaurant choice (Banerjee, 1992). The state of nature is such that, of two available restaurants, *Left* is better than *Right.* These two restaurants provide the options between which the agents must choose, and each agent prefers to go to *Left* if *Left* is the better restaurant, and each prefers to go to *Right* if *Right* is the better restaurant.

The decision is made under uncertainty, in the sense that no agent knows the state of nature, though it is common knowledge that everybody prefers to go to the better restaurant. Specifically, the following information is available:

1. the underlying structure, including the sequence in which agents make their choices, is common knowledge,
2. it is common knowledge that each agent makes a rational decision based on their 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.

In the example, agents may have read a review from home, indicating that *Left* is better than *Right*, or heard from a friend that *Right* is better than *Left*. The other agents do not have access to this private information: 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: everybody can see what everybody before them saw.

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.

3.2.2 Initiating a Cascade

Let our set of agents, call them {*a,b,c,...*}, stand in line in alphabetic order, waiting to make their choice, each informed by their private signals. The choice for the first agent, *a*, is easy: say she received a signal, *left*, indicating that *Left* is the better restaurant. In that case, she will rationally choose the left restaurant.

The second agent, *b*, sees this choice, and from it, he may infer what signal *a* received: if *a* had received signal *right*, then she would have chosen *Right*. She didn't; hence she got signal *left*. Assume now that *b* also received signal *left*. His choice is as easy as *a*'s: he only has reason to choose *Left*.

The third agent, *c*, has seen the choices of both *a* and *b*, and may reason as *b* regarding *a*'s signal: *c* knows *a* received a *left* signal. What does *c* know about *b*'s signal? Does *c* consider it possible that *b* received a *right* signal, but chose *Left*? There is a subtlety here, that requires an assumption about the agents' *tie-breaking rule*. We make the additional assumption that it is common knowledge that if an agent has equally many signals indicating each restaurant, then the agent will choose in accordance with her own private signal. Under this assumption, *c* may conclude that *b* received a *left* signal – if *b* had received a *right* signal, then *b* would know of both a *left* and a *right* signal, and therefore follow his own private signal, *right*.

Hence *c* knows that two *left* signals have been given. If *c* also received a *left* signal, she should clearly chose *Left*. But what if she received a *right* signal? Well, all signals are known to be more likely to be correct than incorrect, so *c* can conclude that it is more likely that the *left* signal is correct, and the *right* signal incorrect. Hence, *c* will, completely rationally, choose *Left* – contrary to her private signal! Hence, *c* is in an information cascade.

The fourth agent, *d*, will also be in cascade. In fact, *d* will be in the same epistemic situation as *c*, as *d* cannot deduce *c*'s private signal. This is a corollary of *c* being in cascade: since *d* knows that *c* is rational and received the public signal (*left*, *left*), *d* can deduce that *c* would have chosen *Left no matter what private signal she received*. Hence, *d* will base his decision only on the choices of *a* and *b*, 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.

3.2.3 Positive and Negative Cascades

In the example above, we arbitrarily specified which restaurant was the better one. Though all agents chose *Left*, the better restaurant could have been *Right*. It is less likely that the initial segment of private signals would have been *left, left* in that case, but not impossible. Moreover, the cascade would still have been rational.

This hints at the strength of cascades: even perfectly rational agents may be caught in a *negative cascade*, a cascade leading to the undesired outcome. It is more likely that agents benefit from using social proof, as *positive cascades* are more likely to occur. However, if the probability of a private signal is correct is 2/3, there is still a 1/9 probability that a negative cascade occurs due to the first two agents! Given that cascades can occur in investment situations and in relation to public opinion before elections, this is not negligible.

3.2.4 Cascades in the Wild

The structure of cascades provided above makes very strong assumptions about the rationality of agents and their available information. In particular, the fact that agents are able to reason indefeasibly about the higher-order information of others – and thereby either deduce their private signals or that they are in a cascade – facilitates the occurrence of cascades but also ensures that such cascades are *fragile*: if agent *e* received *two* signals indicating *Right*, she would be able to break the cascade (as the signals from *c* and *d* cannot be deduced*)*.

One striking assumption of the model is that the *social network* is common knowledge, and that it is furthermore known exactly how information travels through it. In real life, we seldom know that the action of *a* directly influenced *b* or that *c* was influenced only by *a* and *b*. More likely, we have no clue about the informational pathways. This entails that we cannot take a hyper-rational approach to social proof, but must rather rely on practical heuristics. Alas, the common heuristic applied is to assume that decisions are made on a *privately* informed basis, not on social proof. This is less than ideal, as it may facilitate stronger and more frequent cascades.

To illustrate the point, assume that Elise is to form an opinion about whether a new film is worthwhile. She has received no private signals, but seeks the advise of Alice, Bob, Carol and Dale, whom she trusts equally. Alice tells her that she heard the film was bad, while Bob, Carol and Dale tell her that they heard the film was good. If Elise assumes each is privately informed, the evidence to go to see the film is swaying. She would even have to read *two* bad reviews before she would be convinced that Bob, Carol and Dale were off the truth track. However, Bob, Carol and Dale may not be *independent information sources*: if Bob told Carol that the film was good, and Carol told Dale, then really these three should count as only *one piece of evidence* for the film being good. Hence, applying the heuristic from above puts Elise in a cascade that more information about her information sources would have prevented.

The example is fictitious, but the structure reoccurs in most of our institutionalized information aggregation systems: just consider that best-seller lists, online download counts, opinion polls, crowd-based opinion aggregators and even academic citations may all be the victims of such "infostorms" (Hendricks & Hansen, 2014)

3.3 Group Polarization

Shoppers on Amazon 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. Google by default uses your past 180 days' search history to provide Personalized Search for Everyone.[2]

A further common feature of modern web technologies is their being *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 threads. 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 re-share it with their social network (Hendricks & Hansen, 2014). Hereby, our private opinions are shaped by social deliberation.

In relation to social deliberation, 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 (Sunstein, 2009). The phenomenon can reliably be reproduced in lab settings ((Myers and Lamm, 1976), (Myers 1982) for reviews of experimental literature), using, among others, a setup like the following.

3.3.1. The Structure of Polarization

Group polarization may 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

---

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

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 favoured. The shift is measured by comparing the average degree of individual pre-discussion expressions of agreement with a post-discussion expression. The latter may be given either by asking for post-discussion 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 cadets, and university students. In the former 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 favoured the softer responses, whereas officers recommended the more militant solutions. After discussion, these two groups polarized, whereas the ROTC cadets where more neutral in both pre-discussion and post-discussion 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 (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 favour of the leaned-to pole, and that more arguments in favour 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. (Bishop and Myers, 1974) have suggested and supported the view that the group shift is based on a number of parameters, namely, the direction of argument (which pole the group favours), 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 favour 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.

It is 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 favour 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.

4.0 Bubble Studies

Bubbles are typically associated with situations in finance in which assets trade at prices far exceeding their fundamental value (Vogel 2010). Stock and real-estate may get overheated but the same goes for opinions on the web, social status and a whole range of other phenomena in science and society. Now opinion, recognition or social capital is the liquidity to be invested in public viewpoints, fame or online-respect and "likes". One may accordingly consider opinion bubbles, political bubbles, bubbles of social capital, bullying bubbles, polarization bubbles, science bubbles etc. (Hendricks 2014a).

It turns out, from an information theoretical perspective, that bubbles may essentially be viewed as *information control problems* (Hansen, Hendricks & Rendsvig 2013), (Hendricks & Hansen 2014) among deliberating agents who are collectively susceptible to robustly demonstrated socio-psychological features like boom-thinking, group-thinking and lemming effects. These, together with determinate market models and conditions, may make for bubble-hospitable environments over disparate ontologies ranging from cyber-bullying on social media (Hendricks 2014b), to research funding and science bubbles (Budtz Pedersen & Hendricks 2013), (Hendricks 2014c). Further bubble examples founded on social proof and particular market-conditions involve buying the same stock (Shiller 03); thinking the same thing; holding the same opinion online (Hendricks & Hansen 14); subscribing to the same political program (McCarty, Poole, Rosenthal 13); converging on the same enemies virtually or for real; all members thinking the same as the chairman of the board (Halbesleben & Buckley 04); appreciating the same art; "liking" the same posts on social media (Centola 10); taking the same medicine. Everybody is trending the same way dependent on the way in which the individual agent processes the available information about other agents' beliefs, norms and actions but independently of whether this mode of operation is necessarily tracking the truth or is the right thing to do – irrational group behavior or wrongful belief aggregation fuel bubbles (Ofek & Richardson 03), (Hansen, Hendricks & Rendsvig 13), (Hendricks & Hansen 2014).

That bubbles over different ontologies may be viewed as information control problems in networks subject to social proof is intrinsically connected to the current information-driven models of bubble emergence in economics. In particular (Abreu & Brunnermeier 03, Brunnermeier 08) isolate four main strands of bubble models including

(i) models in which all investors have rational expectations and symmetric information (Blanchard & Watson 82),
(ii) models for which investors are asymmetrically informed and the presence of a bubble is not common knowledge (Allen, Morris & Postelwaite 93), (Brunnermeier 13),
(iii) models where bubbles persist due to limited arbitrage because rational and well-informed investors interact with noise traders psychologically biased in unfortunate ways (DeLong et al. 90) and finally
(iv) models of bubbles in which different investors hold different beliefs about the fundamental value of the asset and agree to disagree accordingly (Harrison & Kreps 78).

In all four model types, social information implicitly plays a key role but no over-arching information theory is *yet* present. Indeed problems of informational interaction, socio-psychological influence and information flows across networks of agents or investors are acknowledged by all parties and all models of bubble formation and so "while we have a much better idea of why rational traders are unable to eradicate the mispricing introduced by behavioral traders, our understanding of behavioral biases and belief distortions is less advanced." (Brunnermeier, 2008:14). This is where formal models of socio-informational phenomena like the ones presented here come in – they are part of uncovering and understanding the structure and dynamics of bubble formation.

Information may indeed be used for enlightenment, insight, education and qualified decision and deliberation, but may unfortunately also be used to manipulate people, opinions and markets. Sometimes the manipulation may be intentional – say for people to acquire certain consumer goods and financial products or subscribe to particular political or religious programs – and sometimes agents are the victims of manipulation as the unintentional result of wrongful collective information processing and erroneous group reasoning.

Thus, one information problem to address is to formulate *intervention strategies* for malignant bubbles like unjustified Twitter-storms (#marius or #voteman) where for instance a false tweet from Associated Press crashed the American stock markets in minutes or got the euro to plunge against the US dollar with a false rumour to the effect that the chairman of the German Bundesbank was about to resign (Hendricks 2014d). Another example relates to the strange bubble economics of selfies where social capital is used to overheat fame (Hendricks 2014e). Also polarization, radicalization and extremism may be considered as unfortunate and destabilizing bubble formation (Hendricks 2014f). Such "infostorms" (Hendricks & Hansen 2014, demonstrate how information technology and social media may amplify irrational group behaviour. In this way, *bubbles refer to unfortunate (irrational) ways of collective aggregating behaviour, opinions, preferences or actions based on social proof and marketplaces in science, society and elsewhere*.

But bubbles may not necessarily all be malignant if they mirror public conviction on correct information and social influence rails reason. In economics rational bubbles may exist in which it is reasonable for investors to continue their investment behaviour all the way to bubble emergence. Could there be benign bubbles calling for crowd climate awareness, race and gender equality, health care benefits, anti-radicalization, anti-echo-chambering of ideologies or religious disagreement etc? How benign bubbles may be *stimulated and used to promote* good ideas and socially desirable initiatives is also a very important information control problem to be addressed.

References

Abreu, D. and Brunnermeier, M. K. (2003). "Bubbles and crashes", *Econometrica* 71: 173– 204.

Allen, F., Morris, S. and Postlewaite, A. (1993). "Finite bubbles with short sale constraints and asymmetric information," *Journal of Economic Theory* 61, 206–29.

Blanchard, O. J. and Watson, M. W. (1982). "Bubbles, rational expectations, and financial markets." in *Crisis in the Economic and Financial Structure*, ed. P. Wachtel. Lexington, MA: Lexington.

Brunnermeier, M.K. (2008). “Bubbles”, *The New Palgrave Dictionary of Economics*, Second Edition, S.N. Durlauf & L.E. Blume (eds.), London: Palgrave.

Brunnermeier, M.K. (2013). *Asset Pricing under Asymmetric Information*. Oxford: Oxford University Press, reprinted from 2001.

Banerjee, A.V. (1992). “A Simple Model of Herd Behavior”, *The Quarterly Journal of Economics*, Vol. 107, No. 3, (Aug., 1992): 797-817.

Bicchieri, C. and Fukui, Y. (1999). ”The great illusion: ignorance, informational cascades, and the persistence of unpopular norms”, *Business Ethics Quarterly* 9(1): 127–155.

Bikhchandani, S., Hirshleifer, D. & Welch, I. (1998). "Learning from the Behavior of Others: Conformity, Fads, and Informational Cascades,"*Journal of Economic Perspectives*, 12: 151-170.

Bishop, G.D. and Myers, D.G., (1974) ”Informational influences in group discussion”. *Organizational Behavior and Human Performance* 12: 92-104.

Bjerring, J.C, Hansen, J.U., Pedersen, N.J.L.L. (2014). “On the rationality of pluralistic ignorance,” *Synthese, 191(11): 2445-2470.*

Budtz Pedersen, D. & Hendricks, V.F. (2014). “Science Bubbles”, *Philosophy and Technology*, 27(2014): 503-518.

Cialdini, R. (2007). *Influence: The Psychology of Persuasion,* HarperCollins Publishers, 2007.

Centola, D., Willer, R. & Macy, M. (2005). "The Emperor's Dilemma: A Computational Model of Self-Enforcing Norms", *American Journal of Sociology*, 110, 4: 1009-1040.

DeLong, J. B., Shleifer, A., Summers, L. H. and Waldmann, R. J. (1990). "Noise trader risk in financial markets," *Journal of Political Economy* 98, 703–38.

Halbesleben, J.R.B. og M.R. Buckley (2004). „Pluralistic Ignorance: Historical Development and Organizational Applications“, *Journal of Management History* 42(1): 126-38.

Hansen, P.G., Hendricks, V.F. & Rendsvig, R.K. (2013). “Infostorms”, *Metaphilosophy*, vol. 44(3), April: 301-326

Harrison, J. M. and Kreps, D. (1978). "Speculative investor behavior in a stock market with heterogeneous expectations." *Quarterly Journal of Economics* 89, 323–36.

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

Hendricks, V.F. & Hansen, P.G. (2014). *Infostorms: How to Take Information Punches and Save Democracy*. New York: Copernicus Books / Springer.
Hendricks, V.F (2014a). ”From the Art World to Fashion to Twitter, We're All Living in Bubbles,” *Epoch Times* 12.01.2014: http://www.theepochtimes.com/n3/445574-from-the-art-world-to-fashion-to-twitter-were-all-living-in-bubbles/)
Hendricks, V.F (2014b). ”If You Really Want To Help A Troubled Teen, Don't Like Their YouTube Video”, *Business Insider* 13.2.2014: http://www.businessinsider.com/if-you-really-want-to-help-a-troubled-teen-dont-like-their-youtube-video-2014-2?IR=T)
Hendricks, V.F (2014c). ” Neuroscience risks being the next science research bubble”, Medicalxpress 05.11.2014: http://medicalxpress.com/news/2014-11-neuroscience-science.html)
Hendricks, V.F (2014d). "When Twitter Storms Cause Financial Panic, New Statesman 22.01.2014: http://www.newstatesman.com/business/2014/01/when-twitter-storms-cause-financial-panic).
Hendricks, V.F (2014e). ”The Strange Bubble Economics of Selfies”, *Mashable* 15.05.2014: http://mashable.com/2014/05/15/social-capital/).
Hendricks, V.F (2014f). ”Denmark must not succumb to polarisation in the wake of Copenhagen attacks, *The Conversation*, 16.02.2015: https://theconversation.com/denmark-must-not-succumb-to-polarisation-in-the-wake-of-copenhagen-attacks-37083)
Isenberg, D.J. (1986) ”Group polarization: A critical review and meta-analysis.”, *Journal of Personality and Social Psychology*, 50(6): 1141-1151.
Katz, D. & Allport, F.H. (1931). *Student Attitudes*. Syracuse, N.Y.: Craftsman.
Krech, D. & Crutchfield, R.S. (1948). *Theories and Problems of Social Psychology*. New York: McGraw-Hill.
Latané, B., Darley, J.M. (1970). *The Unresponsive Bystander: Why Doesn't He Help?* New York: Appleten-Century-Crofts.
Latané, B., Nida, S. (1981). “Ten Years of Research on Group Size and Helping”, *Psychological Bulletin*, vol. 89(2): 308-324.
McCarty, N., Poole, K.T. & Rosenthal, H. (2013). *Political Bubbles*. Princeton: Princeton University Press.
Miller, D.T. and McFarland, C. (1987): “Pluralisitic Ignorance: When similarity is Interpreted as dissimilarity”, *Journal of* Personality and Social Psychology, 53: 248-305.
Myers, D.G. and Lamm, H. (1976) “The group polarization phenomenon”, *Psychological Bulletin*, 83(4): 602-627.
Myers, D.G. (1982). “Polarizing Effect of Social Interaction” in Brandstätter, H., Davis, J.H. and Stocker-Kreichgauger, G. (eds.) *Group Decision Making*. Academic Press.
Ofek, E. and Richardson, M. (2003). "DotCom mania: The rise and fall of Internet stocks." Working Paper No. FIN-01-037 58(3), 1113–38. New York University, Stern School.
Prentice, D.A, Miller, D.T. (1993). “Pluralistic ignorance and alcohol use on campus: Some consequences of misperceiving the social norm,” *Journal of Personality and Social Psychology*, 64(2): 243–256.
Rendsvig, R.K. (2014. “Pluralistic ignorance in the bystander effect: informational dynamics of unresponsive witnesses in situations calling for intervention”. *Synthese*, vol. 191(11), 2471-2498.
Shiller, R. (2004). *Irrational Exuberance*. Princeton: Princeton University Press.

Sunstein, C.R. (2009). *Going to Extremes: How Like Minds Unite and Divide*. Oxford: Oxford University Press.

Surowiecki, J. (2004). *The Wisdom of Crowds: Why the Many Are Smarter Than the Few and How Collective Wisdom Shapes Business, Economies, Societies and Nations*. New York: Doublday

van Benthem, J. (2014). *Logical Dynamics of Information and Interaction*. Cambridge: Cambridge University Press.

Vogel, H.L. (2010). *Financial Markets: Bubbles and Crashes*. New York: Cambridge University Press.

Westphal, J.D. and Bednar, M.K. (2005). “Pluralistic Ignorance in Corporate Borads and Firms' Strategic Persistence in Response to Low Firm Performance”, *Administrative Science Quarterly* (50): 262-298.