Cognitive and Neural Mechanisms of Social Influence in Decision Making

Inauguraldissertation
zur
Erlangung der Würde
eines Doktors der Philosophie
vorgelegt der
Fakultät für Psychologie
der Universität Basel

von

Rafael Erich Huber

aus Stallikon, Zürich

Basel, 2014

UNI
BASEL

Genehmigt von der Fakultät für Psychologie

auf Antrag von

Prof. Dr. Jörg Rieskamp
Prof. Dr. Rainer Greifeneder

Basel, den ________________________

____
Prof. Dr. Roselind Lieb

UNI
BASEL

Cognitive and Neural Mechanisms of Social Influence in Decision Making

Rafael E. Huber

Department of Psychology

University of Basel

Brief dissertation framework for the cumulative dissertation based on:

Huber R. E., Klucharev V., & Rieskamp J. (2014)

Huber R. E., Herzog S. M., Horn S. S., Klucharev V., & Rieskamp J. (2014)

Schöbel M., Rieskamp J., & Huber R. E. (2014)

Submitted by: Rafael E. Huber (University of Basel)

Advisors: Vasily Klucharev (National Research University - Higher School of Economics, Moscow)

Jörg Rieskamp (University of Basel)

Reviewers: Jörg Rieskamp (University of Basel)

Rainer Greifeneder (University of Basel)

Chair: Jens Gaab (University of Basel)

Declaration

I, Rafael E. Huber (born May 4th 1983 in Zürich, Switzerland), hereby declare the following:

(i) My cumulative dissertation is based on three manuscripts (Huber, Herzog, Horn, Klucharev, & Rieskamp, 2014; Huber, Klucharev, & Rieskamp, 2014; Schöbel, Rieskamp, & Huber, 2014). I contributed substantially and independently to all three manuscripts in this dissertation. The specific contributions are:

- Huber, Herzog et al. (2014): Jointly responsible for the idea and development of the paradigm. Primarily responsible for the data collection, all the analyses (except the estimation of the drift diffusion model parameters) and the writing of the paper.
- Huber, Klucharev, and Rieskamp (2014): Jointly responsible for the idea and development of the adapted version of the paradigm. Primarily responsible for the data collection, all the analyses and the writing of the paper.
- Schöbel et al. (2014): Primarily responsible for the computational modeling and partly responsible for writing the paper.

(ii) I only used the resources indicated.

(iii) I marked all the citations.

Basel, 8th of May, 2014

Rafael E. Huber

Acknowledgements

My dissertation is “social” through and through. While doing research on social influence in decision making, I received tremendous support from my own social environment. First and foremost I would like to thank my family. My father, my mother, my brother and my grandmother always backed me up to the fullest of their possibilities: Thank you! I will always love you. I would also like to thank my close friends with whom I shared so many unforgettable moments. You helped me to keep the bigger picture of life in mind and to ground myself during those times when science became a little bit too important. I cannot name you all, but I know that you know who you are. A very big thank you goes to my supervisor Vasily Klucharev. It was inspiring and motivating to work with someone who truly loves what he is doing. All the best for your future endeavors, Vasily! I am also very much indebted to my Ph.D. advisor Jörg Rieskamp. Without having the opportunity to profit from such a tremendous knowledge in computational modeling, statistics and decision science this dissertation would have never been possible. Stefan Thommen, Laurent Muller, Oliver Schürmann and all the other junior research assistants who supported my research through the years: You hard-working, smart and dependable guys did an invaluable job! Last but not least I would like to thank Laura Wiles for editing several of my manuscripts, my co-authors Markus Schöbel, Stefan Herzog and Sebastian Horn for the pleasant collaboration, Rainer Greifender and Jens Gaab for being part of the examination committee, my former students for all the instructional experiences, and all present and past members of the Economic Psychology lab with whom I spent a wonderful time in Basel. Thank you all!

Abstract

Often insufficient information creates a situation in which we are forced to decide under uncertainty. In such a situation the behavior of others can complement private information and decisively influence a final decision. In many cases relying on the behavior of others is a good strategy and results in more accurate decisions. However, from time to time the information derived from the behavior of others is wrong and relying on such misleading information can trigger herds with destructive consequences (e.g., on the stock market). To better understand how herding behavior develops, methods from computational modeling and neuroscience were combined with theories from social psychology and economics. In the first manuscript a straightforward categorization task was analyzed with a prominent computational model to describe how opinions from others can influence the decision process. That people often treat private information in a privileged way is shown in a second manuscript. It suggests a neural mechanism on how overweighting of private information changes belief updating. Understanding this process is important, as overweighting of private information can decrease the probability that herds develop. Importantly, if private information is overweighted strongly depends on the type of social information, which is shown in a third and final manuscript. The analyses demonstrate that private information is only overweighted as compared to social information derived from the decisions of equally ranked others, but not as compared to social information derived from higher ranked others. In sum, this dissertation sheds light on social influence and the development of herding behavior by studying individual decisions on the psychological and neural level of implementation. Even herding behavior is a group phenomenon it ultimately rests on the wrong decisions of individuals. A better understanding of the associated mechanisms is crucial for the understanding of how fatal herds, as the ones on the stock market, can develop.

Cognitive and Neural Mechanisms of Social Influence in Decision Making

Approximately seven years after its onset and rapid development into a worldwide financial crisis, the implications of the U.S. subprime mortgage crisis are still having a destructive effect on the lives of many people. According to Robert J. Shiller, winner of the 2013 Nobel Prize in Economics, the U.S. subprime crisis is based on a speculative bubble in the housing market that broke in 2006 (Shiller, 2008). Bubbles in financial markets are an emergent phenomenon on the macro-level, but they ultimately rest on the wrong (and misleading) decisions of individual agents on the micro-level. Such individual decisions can be strongly influenced by the opinions and decisions of other agents (Cialdini & Goldstein, 2004). Even the conventional stock valuation theory assumes "[…] that a stock's current market value tends to converge to the (risk adjusted) discounted present value of the rationally expected dividend stream." (Smith, Suchanek, & Williams, 1988), bubbles impressively demonstrate that people's expectations can strongly deviate from this rational prospect. A better understanding of the causes that lead to such devastating consequences of financial bubbles is a highly interdisciplinary endeavor. It requires the combination of knowledge from various scientific disciplines such as economics, sociology, psychology, evolutionary biology and neuroeconomics (Baddeley, 2010). The overarching goal of this dissertation was to deepen our knowledge of how the mechanisms of decision making are influenced by various social factors (e.g., authority or group size). Understanding these mechanisms is essential because they can explain how the wrong and misleading decisions that cause financial bubbles arise. Three manuscripts report the mechanisms of decision making that were studied on the level of cognitive and neural information processing by combining theories from social psychology and economics with tools, models and techniques from statistics, cognitive psychology and neuroeconomics. In a first manuscript (Huber, Herzog, Horn, Klucharev, & Rieskamp, 2014), we report how the cognitive mechanism of social influence (Germar, Schlemmer, Krug, Voss, & Mojzisch, 2013) is modulated by an increase in the size of a group. We found that the effect of an increase in group size from one to 19 on the propensity to conform goes along with a more efficient processing of sensory information. This perceptual bias towards the choice option favored by others is accompanied by a group size dependent increase in response cautiousness. These results confirm earlier findings by Germar et al. (2013) and bring research on the functional relationship between group size and conformity (Bond, 2005) to a round figure. In a second manuscript (Huber, Klucharev, & Rieskamp, 2014), we studied how a bias towards private as compared to social information modulates belief updating. This question is important for a better understanding of herding behavior because the probability that a cascade will start decreases when people put too much weight on their own private information (Nöth & Weber, 2003). Our main findings suggest that the more people overweight private information, the more activity can be observed in the inferior frontal gyrus/anterior insula and the less activity can be observed in the parietal-temporal cortex when people update their beliefs by private information. These results on the neural level point to a two-fold psychological mechanism with emotional and cognitive risk-processing components (Loewenstein, Weber, Hsee, & Welch, 2001). A third and final manuscript (Schöbel, Rieskamp, & Huber, 2014) reports how we used computational modeling to study how a change in the social environment affects information weighting in situations prone to herding behavior. In a first experiment, we replicated the classic urn and balls study on rational herding by Anderson and Holt (1997). The main conclusion is that people have a general tendency to overweight their own private information. In a second experiment, we transferred the abstract urn and balls setting to an ecologically more valid environment. We observed that people who made decisions in this more realistic setting overweighted decisions from higher ranked individuals as compared to decisions from equally ranked peers. Weighting of social information seems to depend strongly on authority information – an often neglected factor in previous studies. As already mentioned, all three manuscripts have the common overarching goal of deepening our knowledge of the cognitive and neural mechanisms underlying social influence in decision making. The remainder of this framework consists of two parts. In a first part, the reader will be provided with a short historical and theoretical overview of the previous work on which this dissertation is built. This first part is essential for understanding the second and final part of the framework: A short summary of all three manuscripts.

Fundamentals of Herding Behavior: Theories From Psychology and Economics

According to Raafat, Chater, and Frith (2009), herding can be “[…] defined as the alignment of the thoughts or behaviors of individuals in a group (herd) through local interaction and without centralized coordination.” The general taxonomy introduced by these authors distinguishes between global, pattern-based (i.e., connections between agents) and local, transmission-based (i.e., exchange of information between agents) mechanisms of herding. Importantly, these two mechanisms almost always work in a highly interconnected way. However, the work described in this dissertation builds more heavily on the idea of transmission-based mechanisms and predominantly on the cognitive (as compared to the affective) aspects of herding behavior. Theories from social psychology as well as rational models (e.g., informational cascades) developed in economics build the core of this branch of herding research. In order to develop paradigms that are optimally suited to studying the cognitive and neural mechanisms of social influence in decision making, we combined the advantages of both approaches. They will be described in the following two paragraphs.

In the field of social psychology the early work of Solomon Asch (1951, 1952, 1955, 1956) is regarded by many as the starting point of conformity research. In his classic experiments on the line judgment task, Asch confronted participants with lines of different length. The seemingly simple task was to decide which of these lines has equal length to an additionally presented reference line. The task was indeed very simple – participants who solved this problem alone in the control condition made almost no mistakes. However, participants in the experimental condition, who solved the problem after several confederates unanimously stated a wrong answer, made mistakes in more than a third of the trials and only about every fourth participant did not make any mistakes at all. Asch recognized that increasing the group size can lead to an increase in the propensity to conform. However, the idea that a majority size of three already exerts the full impact is nowadays questioned (Bond (2005) provides an excellent review of this topic).

Initial research on conformity had a tendency to highlight the negative side of social influence (Larrick, Mannes, & Soll, 2012). According to this view, people put too much weight on the information provided by others. Importantly, Morton Deutsch and Harold B. Gerard (1955) pointed out that one has to distinguish between informational and normative social influence. Informational social influence refers to people's motivation to gain a more accurate perception of reality, whereas normative social influence refers to people's motivation to be an accepted member of a group. Here, relying on other people's informational social influence can often improve decisions (Larrick, Mannes, & Soll, 2012; Mannes, 2009; Surowiecki, 2005), whereas only following others because of their normative social influence can be both good and bad, depending on the situation. Although initial pioneers in the field of conformity research did not distinguish between these two concepts, newer research tends to focus more strongly on the advantages of following others. Interestingly, researchers studying advice-taking (Bonaccio & Dalal, 2006; Yaniv, 2004; Yaniv & Kleinberger, 2000) found that – contrary to the classic view – people sometimes even put too low a weight on the opinions of others (*egocentric advice discounting*). This view on conformity is partly supported by a second branch of research that originated in economics – the research on informational cascades (Anderson & Holt, 1997; Banerjee, 1992; Bikhchandani, Hirshleifer, & Welch, 1992).

About forty years after the initial work in the field of social psychology, researchers in the field of economics started to study herding behavior with a radically different approach. Theories following this approach are built on the assumptions of rational expectations theory;

that is, they assume independent, rational and self-interested agents who use all the available information without making systematic mistakes (Baddeley, 2010). Such rational herding models have been widely used to explain people's behavior in the lab (Anderson & Holt, 1997; Hung & Plott, 2001), in the labor market (Oberholzer-Gee, 2008), and in financial markets (Chari & Kehoe, 2004; Devenow & Welch, 1996). Probably the most prominent among these rational herding models is the theory of informational cascades. Informational cascades demonstrate that when people decide sequentially without revealing their private information, situations can occur in which following the precedent others and deciding against one's own private information can be the best one can do. If people follow the underlying assumption of the model to a sufficient degree (that is, if they update their beliefs in a way which is consistent with the normative solution provided by Bayes), herds will occur even though people did the best they could have done in a particular situation. However, research has shown that people often tend to overweight their own private information, even in situations in which it would be best for them to follow others (Weizsäcker, 2010). Importantly, such overconfident overweighting of private information can decrease the probability that cascades start and/or persist (Bernardo & Welch, 2001; Nöth & Weber, 2003). A major goal of this dissertation is a better understanding of how the cognitive process of belief updating is influenced by changes in the environment (Schöbel et al., 2014) and how a bias towards private information modulates the neural mechanism underlying belief updating (Huber, Klucharev, & Rieskamp, 2014).

The outlined research in social psychology and economics has shown that (a) people's decisions are influenced by others (in a good and in a bad way), (b) this influence can depend on the specific characteristics of the environment (normative vs. informational social influence, group size, …), and (c) people in certain situations tend to overweight their own private as compared to the available social information. Understanding the cognitive and neural mechanisms of social influence in decision making requires some knowledge of how these mechanisms work in general that is without social influence. The next section will therefore provide the reader with some fundamentals on the neural mechanisms of belief updating and decision making under uncertainty. In a subsequent section, a short introduction to the drift-diffusion model (Ratcliff, 1978) exemplifies how mathematical models of decision making can help to disentangle the cognitive mechanisms underlying social influence.

The Neural Basis of Belief Updating and Decision Making Under Uncertainty

When faced with the question of which of several financial products she should invest her money in, a real-world decision maker most often cannot relate these choice options to exact success probabilities. Early 20th century economist Frank H. Knight (1921) introduced the now famous distinction between risk and uncertainty. In situations in which an agent decides under risk, outcome probabilities (and outcomes) are known (that is, they can be logically deduced or inferred from data), whereas in situations of uncertainty information on outcome probabilities is not available (Meder, Le Lec, & Osman, 2013). Ambiguity refers to situations in which both – probabilities and outcomes – can be uncertain and Daniel Ellsberg (1961) famously demonstrated that people generally are ambiguity averse. Meder et al. (2013) pointed out that the differentiation between the two concepts *risk* and *uncertainty* can be problematic. It is often very difficult to qualitatively discriminate between mechanisms of decision making under risk and uncertainty. Therefore, in this dissertation the term uncertainty refers to all forms of uncertainty (including risk), especially in the context of the neural underpinnings of uncertainty.

Several brain structures – among others, the dorsomedial prefrontal cortex (DMPFC), the anterior insula, the dorsolateral prefrontal cortex (DLPFC) and the parietal cortex – have been associated with the neural mechanism underlying decision making under uncertainty (Bach & Dolan, 2012; Mohr, Biele, & Heekeren, 2010; Platt & Huettel, 2008). The different brain structures of this network were associated with different sub processes of decision making under uncertainty. According to a recent meta-analysis by Mohr et al. (2010), the anterior insula (together with the thalamus) is thought to be part of an emotional circuitry. The DMPFC, on the other hand, seems to be involved in more cognitive aspects. Finally, together with the parietal cortex, the DLPFC (Huettel, 2006; Huettel, Song, & McCarthy, 2005) is thought to be important for the process of forming a decision and selecting an action. Interestingly, this last finding was confirmed by a study of Stern, Gonzalez, Welsh, and Taylor (2010), which also found fronto-parietal brain structures to be active while participants executed a decision. Additionally, this paper described a neural mechanism of belief updating by showing that the activity in the dorsal anterior cingulate is related to objective uncertainty while participants accumulate evidence. Another study (D'Acremont, Schultz, & Bossaerts, 2013) distinguished the process of evidence accumulation (objective frequencies) from the process of tracking of Bayesian posterior probabilities (objective frequencies in combination with prior information). Here, evidence accumulation was found to be associated with activity in angular gyri, posterior cingulate and medial prefrontal cortex, whereas tracking of Bayesian posterior probabilities was related to activity in bilateral inferior frontal gyrus. The two studies just described, demonstrate in an exemplary way that the neural mechanism of belief updating is not yet as well understood as the more general neural mechanism of decision making under uncertainty. Huber, Klucharev, & Rieskamp, 2014 provide evidence for the idea that specific parts of the neural network of decision making under uncertainty are modulated by the weight people give to their own private as compared to social information while they update their beliefs. As information integration in cascade situations is ultimately nothing other than belief updating in a social environment, a better knowledge of these mechanisms is crucial for better understanding the causes that lead to informational cascades.

In order to study such complex psychological (and neural) processes as those involved in sequential decision making (Gluth, Rieskamp, & Büchel, 2012) or belief updating (D'Acremont et al., 2013; Stern et al., 2010), the application of computational models has become increasingly popular. The goal of these computational models is to translate a complex and noisy data set (often with several dependent variables) into a set of psychologically interpretable parameters. Two prominent examples, which were also applied in the field of neuroeconomics, are reinforcement learning models (O'Doherty, Dayan, Friston, Critchley, & Dolan, 2003; O'Doherty et al., 2004) and random walk or diffusion process models (Gold & Shadlen, 2007; Mulder, Wagenmakers, Ratcliff, Boekel, & Forstmann, 2012; Philiastides, Auksztulewicz, Heekeren, & Blankenburg, 2011; Philiastides, Ratcliff, & Sajda, 2006). The advantage of computational models for psychology in general and for the study of social influence in particular will be highlighted in the next section. This section especially highlights the diffusion model because, as described in the first manuscript of this dissertation (Huber, Herzog et al., 2014), this model was used to study how the cognitive mechanism of social influence (Germar et al., 2013) is modulated by an increase in group size.

The Drift-diffusion Model – Advantages of Model Based Social Science

More than thirty years have passed since Roger Ratcliff introduced the model known as drift-diffusion model (1978). The diffusion model has been widely used to study two-alternative forced choice tasks in the field of cognitive psychology (see Voss, Nagler, & Lerche, 2013 for a good introduction) and more recently in the field of social psychology (Germar et al., 2013; Klauer, Voss, Schmitz, & Teige-Mocigemba, 2007; Philiastides & Ratcliff, 2013). Interestingly, it took quite some time before the model became more generally accepted. The exponential change in citations of "Rafcliff (1978)" since approximately the mid-nineties nicely reflects this increased interest in the diffusion model (Voss et al., 2013). The very general idea of cognitive models is to build a bridge between behavioral (and/or neural) data and latent psychological processes (Forstmann, Wagenmakers, Eichele, Brown, & Serences, 2011). Cognitive modeling is a very powerful tool. It forces researchers to translate verbal hypotheses into mathematical equations and to make quantitative predictions and comparisons of these different hypotheses (Farrell & Lewandowsky, 2010). Most often cognitive models have not only fixed parameters, but also free parameters, which are estimated by fitting a model to data. This can be performed using traditional statistics (e.g., maximum likelihood or least squares) or – as has been done several times in this dissertation (Huber, Klucharev, & Rieskamp, 2014; Schöbel et al., 2014) – by using Bayesian statistics (Kruschke, 2010a, 2010b, 2011). The drift-diffusion model has seven free parameters (Ratcliff, 1978; Ratcliff & McKoon, 2008; Voss et al., 2013; Wagenmakers, 2009), but often researchers focus on only four of these seven. These four parameters can be clearly interpreted in terms of latent cognitive processes, which has been empirically shown several times (Ratcliff, 2002; Voss, Rothermund, & Voss, 2004). The drift rate is higher the more easily a stimuli can be encoded and is therefore affected by task difficulty. The (relative) starting point reflects an a priori bias towards one of the two decision options. The boundary separation can be understood as response cautiousness – the higher the boundary separation the more evidence a participant needs in order to make a decision. This parameter – together with the (relative) starting point – is thought to be under subjective control of an individual (Wagenmakers, 2009). Last, but not least, the non-decision time parameter reflects the time needed for processes, such as motor preparation, that are not part of the actual decision. The drift-diffusion model takes into account all relevant data (that is, the response time distributions for correct and wrong decisions as well as accuracies) and transforms them into psychologically interpretable parameters. In recent years, this approach has become increasingly popular for studying social phenomenon, e.g., to gain a better understanding of the underlying cognitive processes of the implicit association test (Klauer et al., 2007), branding (Philiastides & Ratcliff, 2013), and social influence (Germar et al., 2013). The first manuscript of this dissertation (Huber, Herzog et al., 2014), which will be described in the next section, shows how the general cognitive mechanism of social influence is modulated by an increase of group size.

Why Does Social Influence Increase With Group Size?

Huber, R. E., Herzog, S. M., Horn, S. S., Klucharev, V., & Rieskamp, J. (2014). Why Does Social Influence Increase With Group Size? A Diffusion Model Analysis.

Since Solomon Asch's famous studies on the line judgment task (1952, 1956) it has become well known that the opinions of others can strongly influence individuals' decisions. Even in his very early work, Asch (1951, 1955) recognized that the size of a group can decisively moderate the effect that others have on individuals' decisions. Since these early days, much research has been carried out and several theories try to functionally relate group size to conformity (see Bond, 2005 for a review and meta-analysis). Most of these theories point to a curvilinear relationship (Asch, 1951; Latané, 1981; Latané & Wolf, 1981; MacCoun, 2012; Mullen, 1983; Tanford & Penrod, 1984), but newer research shows that sometimes a linear relationship can do the job just as well, especially for increases in group size above two (Bond, 2005). Although the functional relationship between group size and conformity has been studied extensively, the same cannot be said about the underlying psychological mechanism. Germar et al. (2013) were the first who successfully applied the diffusion model in order to show that social influence mainly affects perceptual bias and response cautiousness. The initially very plausible alternative hypothesis of a change in judgmental bias was not supported by their data. This is an important finding: The core mechanism of social influence in perceptual decision making seems to be a change in a subjectively uncontrollable mechanism (the ease of encoding). Although it should be mentioned that people also required more evidence before they made a decision – a factor that is thought to be under subjective control (Wagenmakers, 2009) – the second important bias parameter (relative starting point) does not seem to be part of the psychological mechanism of social influence. Although Germar et al.'s (2013) study was able to show a convincing psychological mechanism of social influence, they did not answer the important question of how an increase in group size could alter this mechanism. To answer this question, we combined the face-versus-car categorization task (Philiastides et al., 2011) with group opinions and applied the diffusion model to further disentangle the psychological mechanism. This approach allowed us to test whether an increase in group size leads to a change in the general mechanism proposed by Germar et al. (2013) or whether group size affects an additional parameter, which is not affected by social influence in general (e.g., the relative starting point). Our results support the conclusion that social influence is mainly due to a change in perceptual bias as well as due to a change in boundary separation and show that an increase in group size mainly leads to a parametrical change in the general mechanism of social influence (Germar et al., 2013). Interestingly, the pattern that we found seems to be somewhat incompatible with a linear model of group size and social influence. An increase in group size from zero to one leads to a comparable increase in the relevant parameter values (drift rate and boundary separation) to an increase in group size from one to 19. If the pattern was linear, an increase in group size from one to 19 should be much larger than an increase from zero to one. However, although these findings can be related to research trying to functionally relate group size and conformity (Bond, 2005), they are exploratory in nature and have to be confirmed with future studies. This study offers a plausible explanation of how group size could influence the psychological mechanism leading to an increase in social influence – knowledge that complements research on the functional relationship between these two variables.

Neural Correlates of Informational Cascades

Huber, R. E., Klucharev, V. & Rieskamp, J. (2014). Neural Correlates of Informational Cascades: Brain Mechanisms of Social Influence on Belief Updating. Manuscript submitted for publication.

In recent years, developments in the global economy have impressively demonstrated which disruptive forces can grow up from financial bubbles. Based on the assumption of rationally acting agents, the theory on informational cascades (Anderson & Holt, 1997; Banerjee, 1992; Bikhchandani et al., 1992) has offered an explanation of how financial bubbles as well as other forms of herds can emerge from a series of correct, but unfortunate, decisions. In such sequential decision problems, people can base their decisions on *social information* deduced from the decisions of preceding others and on *private information* that is known only to them. Here, situations can occur in which an individual is confronted with prior evidence from social information that is more convincing and contrary to the evidence provided by the individual's private information. In these so-called informational cascades, individuals are thought to integrate social and private information according to a process of Bayesian belief updating. A decision maker who acts as suggested by the normative Bayesian solution weights each piece of evidence equally before deciding on a final choice option. Importantly, previous studies have shown that people do not always act in harmony with this solution, but sometimes deviate by integrating the available information with a bias towards their own private signal[1] (Bernardo & Welch, 2001; Nöth & Weber, 2003; Weizsäcker, 2010). Individuals who update their beliefs in an "overconfident", biased way increase the probability that cascades do not occur in the first place or terminate prematurely.

In the study presented here, participants in a hypothetical decision scenario acted as stock market traders who have to repeatedly decide which of two stocks they want to buy.

---

[1] Note, that in indifferent situations (i.e., in situations in which social information provides the same amount of evidence as private information, but for the opposite choice option) it can be rational to give a slightly higher weight to one's own private information, if there is a probability > 0 that one of the previous decision makers decided wrongly (Anderson & Holt, 1997).

This paradigm can be seen as an adapted version of the classic informational cascades paradigm (Anderson & Holt, 1997). In combination with fMRI and computational modeling this paradigm enabled us to study the cognitive and neural underpinnings of (biased) belief updating in a social environment; that is, the (psychological) mechanisms on which the development of informational cascades finally rests. Before participants decided which of the two stocks they wanted to buy and provided a final (success) probability judgment, they were sequentially confronted with two decisions of previous traders and an own private recommendation from a rating agency. There are three basic hypotheses: (a) all three pieces of information are weighted differently, (b) social and private information are weighted differently, and (c) all the available information is weighted according to the assumptions of the normative Bayesian solution. All three hypotheses were translated to computational models and compared on a behavioral level.

The behavioral analyses show that participants' choices were, in the vast majority of cases, compatible with the normative Bayesian solution. This, however, is not the case for the probability judgments, where participants' behavior revealed a general tendency towards overweighting of private signals (corresponds to hypothesis b). Interestingly, the more people overweighted private information the less they started a cascade in situations specifically prone to herding. On the neural level, studies conducted in decision neuroscience have convincingly shown that brain structures such as the inferior frontal gyrus/anterior insula, the dorsomedial prefrontal cortex (DMPFC), the dorsolateral prefrontal cortex (DLPFC), and the parietal cortex (among others) are involved in the processes of belief updating and decision execution in decision making under uncertainty (D'Acremont et al., 2013; Mohr et al., 2010; Stern et al., 2010). However, this study goes a step further by postulating a potential mechanism of how the neural underpinnings of belief updating are modulated by a bias towards private information. The more uncertain participants became as a result of belief updating by private information the more brain activity was observed in the DLPFC, DMPFC, inferior frontal gyrus/anterior insula and in the parietal-temporal cortex. Importantly, this process seems to be modulated by how participants weight private as compared to social information: The more participants overweighted private information, the more activity was observed in inferior frontal gyrus/anterior insula and the less activity was observed in the parietal-temporal cortex during belief updating by private information. All in all, this study suggests a neural mechanism underlying biased belief updating – the process that can decisively modulate the probability that cascades occur.

Social Influences in Sequential Decision Making

Schöbel, M.,. Rieskamp, J, & Huber, R. E. (2014). Social Influences in Sequential Decision Making. Manuscript submitted for publication.

According to a prominent theory in social psychology, conformity influences people's behavior via two different routes: informational and normative social influence (Deutsch & Gerard, 1955). This dual-process perspective separates people's motivation to gain a valid and accurate perception of reality (*informational social influence*) from people's motivation to act in accordance with the positive expectations of others (*normative social influence*). Normative social influence is thought to be stronger in tasks in which people have to respond in public, whereas informational social influence seems to predominate when people can provide their answers privately (Bond, 2005). However, social expectations can influence people's behavior in private task settings as well (Wood, 2000). Therefore, the distinction between internalization of beliefs as compared to (public) compliance (Festinger, 1953; Moscovici, 1980) cannot easily be simplified to private versus public task settings. Even most conformity researchers seem to agree that (at least) two types of conformity processes exist, it has been difficult to experimentally separate and quantitatively distinguish these two forces. How these two processes interact is a major question in the field of social psychology (Allen, 1965; Levine & Russo, 1987; Tajfel, 1969).

The theory of informational cascades (Anderson & Holt, 1997; Banerjee, 1992; Bikhchandani et al., 1992) describes how people integrate private and publicly available social information. According to this theory, people update their beliefs in line with the assumptions of rational expectations theory; that is, they act in a Bayesian optimal way (Baddeley, 2010). Although this theory has been successfully applied to describe the occurrence of herds, e.g., in financial markets (Chari & Kehoe, 2004), it has mostly neglected the existence of different social influence processes. In this study, we combined the dual-process concept of conformity and the paradigm of informational cascades with computational modeling in order to quantitatively disentangle informational and normative social influence.

In a first study (Experiment 1), we replicated the general findings of Anderson and Holt (1997) and showed that people act in accordance with the informational influence hypothesis, which states that people will base their decisions on both private and social information. The results from our social influence model show that people have a tendency to overweight their own private as compared to the publicly available social information. This violation of rational expectations theory can decrease the probability that a cascade occurs and has already been discussed previously (Bernardo & Welch, 2001; Nöth & Weber, 2003; Weizsäcker, 2010).

In a second study (Experiment 2), participants solved the abstract informational cascades paradigm in an ecologically more realistic medical decision making context. Here, participants acting as assistant physicians had to decide which of two different diseases a patient suffers from. Both diseases were associated with the same symptoms, but they occurred with a different likelihood. By introducing two types of social opinion sources – that is (a) hierarchically higher ranked medical directors and (b) hierarchically equally ranked assistant physicians – we were able to test a form of normative social influence that is based on authority (Milgram, 1974). The results show that decisions in favor of private information as well as confidence ratings (i.e., probability judgments) were consistently lower in the authority condition as compared to the baseline condition when the authority opinion was contrary to private information. In indifference situations, 61.5% of all participants decided against their private information when confronted with opposing social information derived from medical directors' opinions. This is in a strong contrast to Experiment 1, where in indifference situations 79.9% of all participants decided according to their private information. These results clearly show that authority influence can have a strong impact on the development of informational cascades. This is also reflected in the results of the social influence model, which shows that private information is only overweighted as compared to social information from equally ranked assistant physicians but not as compared to social information from higher ranked medical directors.

We can conclude that (a) it is possible to quantitatively separate normative from informational social influence, (b) it is important for the theory of informational cascades to incorporate different sources of social influence, and (c) it is recommended to focus not only on the individual but also on the environment in which individuals decide, when the goal is to improve the (group) outcome in sequential decision making problems.

General Discussion

Herding behavior on the group level is always based on wrong decisions on the individual level. As a consequence, we can only fully understand herding behavior (e.g., in stock markets), if we know how such wrong decisions are psychologically implemented on the level of the individual. In three manuscripts (Huber, Herzog et al., 2014; Huber, Klucharev, & Rieskamp, 2014; Schöbel et al., 2014) mathematical tools from computational modeling and neuroscience were combined with theories and paradigms from social psychology and economics to study the psychological mechanisms of social influence in decision making. There are three main conclusions: (1) Stronger social influence, due to an increase in group size, leads to a stronger bias in sensory information uptake and to an increase in response caution (Huber, Herzog et al., 2014). (2) Others can create a social environment in which we overweight private as compared to social information and this effect seems to be accompanied by specific changes in the neural network of belief updating (Huber, Klucharev, & Rieskamp, 2014). (3) How people weight private as compared to social information depends on the specific characteristics of the social environment and authority increases the weight assigned to social information.

The first manuscript (Huber, Herzog et al., 2014) demonstrates that social influence in perceptual decision making can result from a bias in sensory information uptake and an increase in response caution. Interestingly, others seem to influence how we "see" the world and this effect is stronger for opinions from 19 others as compared to the opinion of a single other. The interpretation that we "see" the world differently could be tested in future studies by using the method described in Huber, Herzog et al. (2014) in combination with brain imaging techniques (e.g., fMRI). On the one hand this would provide us with deeper insights on how social influence modulates the mechanisms of decision making and on the other hand this could also further validate the psychological interpretation of the parameters of the drift-diffusion model. Further, it would be interesting to know, if the psychological mechanism described in Huber, Herzog et al. (2014) can also explain social influence in situations in which there is no a priori defined correct choice option (e.g., preference for a politician, music star or food item) or in situations in which choice options and their outcomes are coupled in a probabilistic (as compared to a deterministic) way.

In Huber, Klucharev, and Rieskamp (2014) we used cognitive modeling and fMRI to study how people integrate social and private information. In the studied environment, we observed that people overweight their own private as compared to the social information and that this overweighting alters the neural mechanism of belief updating. It is important to better understand how and why people overweight private information because overweighting of private information decreases the probability that herding in the form of informational cascades occurs. The results on the neural level point to a dual-process mechanism with emotional and cognitive components. Here, it would be important to replicate these findings (also with other methods) to provide further evidence for the postulated neural mechanism. Additional knowledge could be gained by changing the social environment (e.g., as in Schöbel et al., 2014) to test, if the same neural mechanism can explain belief updating in different social contexts.

In a third and last manuscript Schöbel et al. (2014) experimentally manipulated the social status of others. This resulted in people giving more weight to social information derived from higher ranked as compared to equally ranked individuals and private information was only overweighted as compared to equally ranked peers. This knowledge should be incorporated in future studies on informational cascades because a change in the weight assigned to different sources of information also changes the probability with which cascades occur. If information of people with a higher social status is also overweighted by decision makers in a professional environment (e.g., medical doctors in a hospital or stock market traders) was not tested and would be a promising idea for future research.

In sum, this dissertation describes different psychological mechanisms of social influence in decision making. These mechanisms can help us understand how the social environment, in which we all decide, can bias important decisions – a process which can result in destructive herds as can be observed again and again in financial markets.

References

Allen, V. L. (1965). Situational factors in conformity. In L. Berkowitz (Ed.), *Advances in Experimental Psychology*, (Vol. 2, pp. 133–175). New York, NY: Academic Press. doi: 10.1016/S0065-2601(08)60105-7

Anderson, L. R., & Holt, C. A. (1997). Information cascades in the laboratory. *The American Economic Review*, *87*(5), 847–862. doi:10.2307/2951328

Asch, S. E. (1951). The effects of group pressure upon the modification and distortion of judgments. In H. Guetzkow (Ed.), *Groups, Leadership and Men: Research in Human Relations* (pp. 177–190). Pittsburgh, PA: Carnegie Press.

Asch, S. E. (1952). *Social psychology*. Englewood Cliffs, NJ: Prentice-Hall.

Asch, S. E. (1955). Opinions and social pressure. *Scientific American*, *193*, 31–35. doi:10.1038/scientificamerican1155-31

Asch, S. E. (1956). Studies of independence and conformity: I. A minority of one against a unanimous majority. *Psychological Monographs: General and Applied*. doi:10.1037/h0093718

Bach, D. R., & Dolan, R. J. (2012). Knowing how much you don't know: a neural organization of uncertainty estimates. *Nature Reviews Neuroscience*, *13*(8), 572–586. doi:10.1038/nrn3289

Baddeley, M. (2010). Herding, social influence and economic decision-making: socio-psychological and neuroscientific analyses. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *365*(1538), 281–290. doi:10.1098/rstb.2009.0169

Banerjee, A. V. (1992). A simple model of herd behavior. *The Quarterly Journal of Economics*, *107*(3), 797–817. doi:10.2307/2118364

Bernardo, A. E., & Welch, I. (2001). On the evolution of overconfidence and entrepreneurs. *Journal of Economics & Management Strategy*, *10*(3), 301–330. doi:10.1111/j.1430-9134.2001.00301.x

Bikhchandani, S., Hirshleifer, D., & Welch, I. (1992). A theory of fads, fashion, custom, and cultural change as informational cascades. *Journal of Political Economy*, *100*(5), 992–1026. doi:10.2307/2138632

Bonaccio, S., & Dalal, R. S. (2006). Advice taking and decision-making: An integrative literature review, and implications for the organizational sciences. *Organizational Behavior and Human Decision Processes*, *101*(2), 127–151. doi:10.1016/j.obhdp.2006.07.001

Bond, R. (2005). Group size and conformity. *Group Processes & Intergroup Relations*, *8*(4), 331–354. doi:10.1177/1368430205056464

Chari, V. V, & Kehoe, P. J. (2004). Financial crises as herds: overturning the critiques. *Journal of Economic Theory*, *119*(1), 128–150. doi:10.1016/S0022-0531(03)00225-4

Cialdini, R. B., & Goldstein, N. J. (2004). Social influence: compliance and conformity. *Annual Review of Psychology*, *55*, 591–621. doi:10.1146/annurev.psych.55.090902.142015

D'Acremont, M., Schultz, W., & Bossaerts, P. (2013). The human brain encodes event frequencies while forming subjective beliefs. *The Journal of Neuroscience*, *33*(26), 10887–10897. doi:10.1523/JNEUROSCI.5829-12.2013

Deutsch, M., & Gerard, H. B. (1955). A study of normative and informational social influences upon individual judgement. *Journal of Abnormal Psychology*, *51*(3), 629–636. doi:10.1037/h0046408

Devenow, A., & Welch, I. (1996). Rational herding in financial economics. *European Economic Review*, *40*(3–5), 603–615. doi:10.1016/0014-2921(95)00073-9

Ellsberg, D. (1961). Risk, ambiguity, and the savage axioms. *The Quarterly Journal of Economics*, *75*(4), 643–669. doi:10.2307/1884324

Farrell, S., & Lewandowsky, S. (2010). Computational models as aids to better reasoning in psychology. *Current Directions in Psychological Science, 19*(5), 329–335. doi:10.1177/0963721410386677

Festinger, L. (1953). An analysis of compliant behavior. In M. Sherif & M. O. Wilson (Eds.), *Group relations at the crossroads* (pp. 232–256). Oxford, England: Harper.

Forstmann, B. U., Wagenmakers, E.-J., Eichele, T., Brown, S., & Serences, J. T. (2011). Reciprocal relations between cognitive neuroscience and formal cognitive models: opposites attract? *Trends in Cognitive Sciences*, *15*(6), 272–279. doi:10.1016/j.tics.2011.04.002

Germar, M., Schlemmer, A., Krug, K., Voss, A., & Mojzisch, A. (2013). Social influence and perceptual decision making: A diffusion model analysis. *Personality and Social Psychology Bulletin*. doi:10.1177/0146167213508985

Gluth, S., Rieskamp, J., & Büchel, C. (2012). Deciding when to decide: time-variant sequential sampling models explain the emergence of value-based decisions in the human brain. *Journal of Neuroscience*, *32*(31), 10686–98. doi:10.1523/JNEUROSCI.0727-12.2012

Gold, J. I., & Shadlen, M. N. (2007). The neural basis of decision making. *Annual Review of Neuroscience*, *30*, 535–574. doi:10.1146/annurev.neuro.29.051605.113038

Huber, R.E., Herzog, S.M., Horn, S.S., Klucharev, V., & Rieskamp, J. (2014). Neural correlates of informational cascades: Why does social influence increase with group size? A diffusion model analysis.

Huber, R.E., Klucharev, V., & Rieskamp, J. (2014). Neural correlates of informational cascades: brain mechanisms of social influence on belief updating. Manuscript submitted for publication.

Huettel, S. A. (2006). Behavioral, but not reward, risk modulates activation of prefrontal, parietal, and insular cortices. *Cognitive, Affective, & Behavioral Neuroscience*, *6*(2), 141–151. doi:10.3758/CABN.6.2.141

Huettel, S. A., Song, A. W., & McCarthy, G. (2005). Decisions under uncertainty: Probabilistic context influences activation of prefrontal and parietal cortices. *Journal of Neuroscience*, *25*(13), 3304–3311. doi:10.1523/JNEUROSCI.5070-04.2005

Hung, A. A., & Plott, C. R. (2001). Information cascades: Replication and an extension to majority rule and conformity-rewarding institutions. *American Economic Review*, *91*(5), 1508–1520. doi:10.1257/aer.91.5.1508

Klauer, K. C., Voss, A., Schmitz, F., & Teige-Mocigemba, S. (2007). Process components of the Implicit Association Test: a diffusion-model analysis. *Journal of Personality and Social Psychology*, *93*(3), 353–368. doi:10.1037/0022-3514.93.3.353

Knight, F. H. (1921). *Risk, Uncertainty, and Profit*. Boston, MA: Houghton Mifflin Company.

Kruschke, J. K. (2010a). Bayesian data analysis. *Wiley Interdisciplinary Reviews: Cognitive Science*, *1*(5), 658–676. doi:10.1002/wcs.72

Kruschke, J. K. (2010b). What to believe: Bayesian methods for data analysis. *Trends in Cognitive Sciences*, *14*(7), 293–300. doi:10.1016/j.tics.2010.05.001

Kruschke, J. K. (2011). *Doing Bayesian Data Analysis: A Tutorial with R and BUGS*. Burlington, MA: Academic Press.

Kübler, D., & Weizsäcker, G. (2003). Information cascades in the labor market. *Journal of Economics*, *80*(3), 211–229. doi: 10.1037/h0046408

Larrick, R. P., Mannes, A. E., & Soll, J. B. (2012). The social psychology of the wisdom of crowds. In J. I. Krueger (Ed.), *Frontiers of Social Psychology: Social Judgment and Decision Making* (pp. 227–242). New York, NY: Psychology Press.

Latané, B. (1981). The psychology of social impact. *American Psychologist*, *36*(4), 343–356. doi:0003-066X/81/3604-0343

Latané, B., & Wolf, S. (1981). The social impact of majorities and minorities. *Psychological Review, 88*(5). doi:10.1037/0033-295X.88.5.438

Levine, J. M., & Russo, E. M. (1987). Majority and minority influence. In C. Hendrick (Ed.), *Group processes. Review of personality and social psychology* (pp. 13–54). Thousand Oaks, CA: Sage Publications, Inc.

Loewenstein, G. F., Weber, E. U., Hsee, C. K., & Welch, N. (2001). Risk as feelings. *Psychological Bulletin*, *127*(2), 267–286. doi:10.1037/0033-2909.127.2.267

MacCoun, R. J. (2012). The burden of social proof: Shared thresholds and social influence. *Psychological Review, 119*(2), 345–372. doi:10.1037/a0027121

Mannes, A. E. (2009). Are we wise about the wisdom of crowds? The use of group judgments in belief revision. *Management Science, 55*(8), 1267–1279. doi:10.1287/mnsc.1090.1031

Meder, B., Le Lec, F., & Osman, M. (2013). Decision making in uncertain times: what can cognitive and decision sciences say about or learn from economic crises? *Trends in Cognitive Sciences*, *17*(6), 257–260. doi:10.1016/j.tics.2013.04.008

Milgram, S. (1974). *Obedience to Authority: An Experimental View*. London, England: Tavistock Publications.

Mohr, P. N. C., Biele, G., & Heekeren, H. R. (2010). Neural processing of risk. *The Journal of Neuroscience*, *30*(19), 6613–6619. doi:10.1523/JNEUROSCI.0003-10.2010

Moscovici, S. (1980). Toward a theory of conversion behavior. In L. Berkowitz (Ed.), *Advances in Experimental Social Psychology* (Vol. 13, pp. 209–239). New York, NY: Academic Press. doi:10.1016/S0065-2601(08)60133-1

Mulder, M. J., Wagenmakers, E.-J., Ratcliff, R., Boekel, W., & Forstmann, B. U. (2012). Bias in the brain: A diffusion model analysis of prior probability and potential payoff. *Journal of Neuroscience*, *32*(7), 2335–2343. doi:10.1523/JNEUROSCI.4156-11.2012

Mullen, B. (1983). Operationalizing the effect of the group on the individual: A self-attention perspective. *Journal of Experimental Social Psychology, 19*(4), 295–322. doi:10.1016/0022-1031(83)90025-2

Nöth, M., & Weber, M. (2003). Information aggregation with random ordering: Cascades and overconfidence. *Economic Journal*, *113*(484), 166–189. doi:10.1111/1468-0297.00091

O'Doherty, J., Dayan, P., Schultz, J., Deichmann, R., Friston, K., & Dolan, R. J. (2004). Dissociable roles of ventral and dorsal striatum in instrumental conditioning. *Science*, *304*(5669), 452–454. doi:10.1126/science.1094285

O'Doherty, J. P., Dayan, P., Friston, K., Critchley, H., & Dolan, R. J. (2003). Temporal difference models and reward-related learning in the human brain. *Neuron*, *38*(2), 329–337. doi:10.1016/S0896-6273(03)00169-7

Oberholzer-Gee, F. (2008). Nonemployment stigma as rational herding: A field experiment. *Journal of Economic Behavior & Organization*, *65*(1), 30–40. doi:10.1016/j.jebo.2004.05.008

Philiastides, M. G., Auksztulewicz, R., Heekeren, H. R., & Blankenburg, F. (2011). Causal role of dorsolateral prefrontal cortex in human perceptual decision making. *Current Biology, 21*(11), 980–983. doi:10.1016/j.cub.2011.04.034

Philiastides, M. G., & Ratcliff, R. (2013). Influence of branding on preference-based decision making. *Psychological Science, 24*(7), 1208–1215. doi:10.1177/0956797612470701

Philiastides, M. G., Ratcliff, R., & Sajda, P. (2006). Neural representation of task difficulty and decision making during perceptual categorization: A timing diagram. *Journal of Neuroscience*, *26*(35), 8965–8975. doi:10.1523/JNEUROSCI.1655-06.2006

Platt, M. L., & Huettel, S. A. (2008). Risky business: the neuroeconomics of decision making under uncertainty. *Nature Neuroscience*, *11*(4), 398–403. doi:10.1038/nn2062

Raafat, R. M., Chater, N., & Frith, C. (2009). Herding in humans. *Trends in Cognitive Sciences*, *13*(10), 420–428. doi:10.1016/j.tics.2009.08.002

Ratcliff, R. (1978). A theory of memory retrieval. *Psychological Review*, *85*(2), 59–108. doi:10.1037/0033-295X.85.2.59

Ratcliff, R. (2002). A diffusion model account of response time and accuracy in a brightness discrimination task: Fitting real data and failing to fit fake but plausible data. *Psychonomic Bulletin & Review*, *9*(2), 278–291. doi:10.3758/BF03196283

Ratcliff, R., & McKoon, G. (2008). The diffusion decision model: Theory and data for two-choice decision tasks. *Neural Computation*, *20*(4), 873–922. doi:10.1162/neco.2008.12-06-420

Schöbel, M., Rieskamp, J., & Huber, R.E. (2014). Social influences in sequential decision making. Manuscript submitted for publication.

Shiller, R. J. (2008). *The Subprime Solution: How Today's Global Financial Crisis Happened, and What to Do about It*. *Crisis*. Princeton, NJ: Princeton University Press.

Smith, V. L., Suchanek, G. L., & Williams, A. W. (1988). Bubbles, crashes, and endogenous expectations in experimental spot asset markets. *Econometrica*, *56*(5), 1119–1151. doi:10.2307/1911361

Stern, E. R., Gonzalez, R., Welsh, R. C., & Taylor, S. F. (2010). Updating beliefs for a decision: Neural correlates of uncertainty and underconfidence. *Journal of Neuroscience*, *30*(23), 8032–8041. doi:10.1523/JNEUROSCI.4729-09.2010

Surowiecki, J. (2005). *The Wisdom of Crowds*. New York, NY: Anchor Books.

Tajfel, H. (1969). Social and cultural factors in perception. In G. Lindzey & E. Aronson (Eds.), *Handbook of social psychology* (2nd ed., Vol. 3, pp. 315–394). Reading, MA: Addison-Wesley.

Tanford, S., & Penrod, S. (1984). Social Influence Model: A formal integration of research on majority and minority influence processes. *Psychological Bulletin, 95*(2), 189–225. doi:10.1037/0033-2909.95.2.189

Voss, A., Nagler, M., & Lerche, V. (2013). Diffusion models in experimental psychology: A practical introduction. *Experimental Psychology*, *60*, 385–402. doi:10.1027/1618-3169/a000218

Voss, A., Rothermund, K., & Voss, J. (2004). Interpreting the parameters of the diffusion model: an empirical validation. *Memory & Cognition*, *32*(7), 1206–1220. doi:10.3758/BF03196893

Wagenmakers, E.-J. (2009). Methodological and empirical developments for the Ratcliff diffusion model of response times and accuracy. *European Journal of Cognitive Psychology*, *21*(5), 641–671. doi:10.1080/09541440802205067

Weizsäcker, G. (2010). Do we follow others when we should? A simple test of rational expectations. *American Economic Review*, *100*(5), 2340–2360. doi:10.1257/aer.100.5.2340

Wood, W. (2000). Attitude Change: Persuasion and Social Influence. *Annual Review of Psychology*, *51*(1), 539–570. doi:10.1146/annurev.psych.51.1.539

Yaniv, I. (2004). Receiving other people's advice: Influence and benefit. *Organizational Behavior and Human Decision Processes*, *93*(1), 1–13. doi:10.1016/j.obhdp.2003.08.002

Yaniv, I., & Kleinberger, E. (2000). Advice taking in decision making: Egocentric discounting and reputation formation. *Organizational Behavior and Human Decision Processes*, *83*(2), 260–281. doi:10.1006/obhd.2000.2909

Why Does Social Influence Increase With Group Size?
A Diffusion Model Analysis

Rafael E. Huber

University of Basel, Department of Psychology, Basel

Stefan M. Herzog and Sebastian S. Horn

Max Planck Institute for Human Development, Center for Adaptive Rationality, Berlin

Vasily Klucharev

National Research University - Higher School of Economics, Department of Psychology, Moscow and University of Basel, Department of Psychology, Basel

Jörg Rieskamp

University of Basel, Department of Psychology, Basel

Author Note

This research was supported by a grant from the Swiss National Science Foundation (100014-130352) to Vasily Klucharev and Jörg Rieskamp. We thank Laurent Muller and Stefan Thommen for their assistance in the data collection and Laura Wiles for editing the manuscript. We also thank Markus Germar and Thom Baguley for helpful comments.
Correspondence concerning this article should be addressed to Rafael E. Huber, University of Basel, Department of Psychology, Missionsstrasse 62A, 4055 Basel, Switzerland, Phone: +41 (0)61 267 02 98, Fax: +41 (0)61 267 06 13, rafael.huber@unibas.ch

Abstract

People's opinions are influenced by the opinions of others. In particular, social influence increases with increases in group size. To understand the underlying cognitive mechanism of this association we study how social influence is modulated by group size using a face-versus-car categorization task. Assuming that people accumulate evidence until a decision threshold is reached this process can be accurately modeled using a drift-diffusion model. Prior to their decisions participants were told the opinion of a small or large group or received no prior information. A large group influenced participants' decisions more than a small group. Modeling the data with the diffusion model revealed that an increase in social influence leads to an increase in the perceptual bias towards a choice option, but at the same time to an increase in response cautiousness. In sum, our cognitive modeling approach illustrates how social influence affects fundamental cognitive judgment processes.

Keywords: social influence, conformity, group size, drift-diffusion model, decision making, judgment processes, perceptual discrimination

Why Does Social Influence Increase With Group Size?

A Diffusion Model Analysis

In his pioneering work on social influence Solomon Asch (1952, 1956) impressively demonstrated how group opinions can influence decision making. Since Asch's initial research the social influence of group opinions on individual judgments has been replicated numerous times (Bond & Smith, 1996). The size of a group is a well-known moderator of this effect (Bond, 2005), with larger groups leading to stronger social influence. Prominent theories examining the association between group size and social influence primarily focused on the mathematical relationship between these two variables (Latané, 1981; Latané & Wolf, 1981; MacCoun, 2012; Mullen, 1983; Tanford & Penrod, 1984). However, how group size affects the underlying cognitive mechanism of the judgment process is not yet fully understood. To overcome this lack of knowledge we examined social influence using an adapted version of the face-versus-car categorization task (Philiastides, Auksztulewicz, Heekeren, & Blankenburg, 2011). In the classic version of the task people have to decide without any further information whether a dynamic noisy visual stimuli depicts a face or a car. In our version of the task people were a priori informed about the opinion of a single other person (1), the opinion of a majority of 19 other persons (2) or they were not informed about the opinion of others (3). To understand the psychological mechanism of increased social influence resulting from an increase in group size we modeled the data using the drift-diffusion model (Ratcliff, 1978).

Group Size and Social Influence

Deutsch and Gerard's (1955) dual-process view separates normative from informational social influence. Whereas normative social influence acts on an individual's desire to be socially approved, informational social influence describes the motivation to get a more accurate perception of reality (Cialdini & Goldstein, 2004). In his meta-analysis on group size and social influence Bond (2005) differentiated between *Asch-type* (face-to-face) and *Crutchfield-type* (individual booths with false group feedback) tasks with private and public response formats. For a *Crutchfield-type* task setting with private responses (including a majority size of one) – comparable to the research presented in this paper – curvilinear models explain the relationship between group size and social influence slightly better than simple linear models. Whereas the initial research by Asch (1951, 1955) as well as the Social Influence model by Tanford and Penrod (1984) assume an asymptotic satiation of social influence (e.g., at a group size of three) other theories, such as the Social Impact theory (Latané, 1981; Latané & Wolf, 1981) or the Other-Total Ratio theory (Mullen, 1983) proposed a negatively accelerated function without an asymptotic limit (Bond, 2005). Recently, MacCoun (2012) introduced the Burden of Social Proof model as a promising addition to existing theories. Different variants of this logistic threshold model can successfully mimic various previous theories.

Interestingly, recent research on the *wisdom of the crowds* effect (Larrick, Mannes, & Soll, 2012; Surowiecki, 2005) provides additional evidence for a curvilinear relationship between group size and (informational) social influence. Integrating an individual opinion with the opinions of others according to a unit-weight strategy that weights both types of opinions equally also results in a negatively accelerated curve (Mannes, 2009). According to a unit-weight strategy an individual opinion and the opinion of one other person both receive an equal weight of .50. When confronted with the opinion of 19 others, the opinion of every person, as well as that of the individual, receives an equal weight of 1/20, which results in a total weight of 19/20 for the group. As a consequence, the increase in the weight assigned to a group opinion, caused by an additional member of a group, strongly diminishes with an increase in group size. Mannes (2009) concluded that, although people tend to (strongly) underweight information provided by large groups, they seem to recognize that larger groups are generally more accurate than small groups or individuals.

In sum, all prominent theories (1) agree that an increase in group size leads to an increase of social influence, (2) mainly focus on the mathematical relationship between group size and social influence and (3) do not provide a detailed account of the underlying psychological mechanism of social influence.

Although the mathematical relationship between group size and social influence has been studied extensively there appears to be a lack of knowledge of how group size affects the judgment process on a psychological level. By using the drift-diffusion model we open this black box and show that the effect of group size on social influence can be explained by a modulation of the general psychological mechanism underlying social influence (see also Germar, Schlemmer, Krug, Voss, & Mojzisch, 2013).

Decomposing Psychological Mechanisms With Sequential Sampling Models

The drift-diffusion model introduced by Ratcliff (1978) belongs to the general class of sequential sampling models. The basic idea of many sequential sampling models is that when presented with a stimuli people start to accumulate evidence for the different choice options until a decision threshold is crossed and a decision is executed (see also Gold & Shadlen, 2007). These models have been applied to a wide variety of cognitive tasks including sensory detection (Smith, 1995), perceptual discrimination (Laming, 1968; Link & Heath, 1975; Usher & McClelland, 2001; Vickers, 1979), categorization (Ashby, 2000; Nosofsky & Palmeri, 1997), probabilistic inference (Wallsten & Barton, 1982), and memory recognition (Ratcliff, 1978). Sequential sampling models have also been successfully applied for value-based decision making (Aschenbrenner, Albert, & Schmalhofer, 1984; Fehr & Rangel, 2011; Gluth, Rieskamp, & Büchel, 2012, 2013a, 2013b; Guo & Holyoak, 2002; Rieskamp, 2008; Roe, Busemeyer, & Townsend, 2001; Usher & McClelland, 2004).

The drift-diffusion model in particular has successfully accounted for behavioral data that is, the shapes of the response time distributions and accuracy from a wide variety of rapid two-choice decision tasks (see Ratcliff & McKoon, 2008; Voss, Nagler, & Lerche, 2013, for reviews). By assuming that evidence from a stimulus is dynamically accumulated over time (starting from a point z) until an internal boundary is crossed, the model disentangles the efficiency of the accumulation process (drift rate parameter v), the amount of information required for the decision (boundary separation a), peripheral nondecision time (T_{er}; e.g., encoding and response execution time), and variability in these components across trials. The drift-diffusion model has also proven to be a very useful tool to better understand the processes involved in social cognition, such as implicit associations (Klauer, Voss, Schmitz, & Teige-Mocigemba, 2007), racial bias (Klauer & Voss, 2008), or the effects of branding (Philiastides & Ratcliff, 2013).

Germar et al. (2013) were the first to use the drift-diffusion model to show that social influence primarily results in people accumulating evidence for the recommended choice option more efficiently (“perceptual bias”). Additionally, people were more cautious (that is, they required more information before they made a decision) when in a situation of social influence. Somewhat surprisingly, social influence did not result in an a priori bias toward a choice option (“judgmental bias”). The findings of Germar et al. (2013) provide first insights into how social influence changes the judgment process. However, how an increase in social influence resulting from an increase in group size affects the judgment process is still unclear and different psychological mechanisms could go along with this modulatory effect of group size. On the one hand, larger group sizes could lead to a linear or non-linear increase in the perceptual bias suggested by Germar et al. (2013). According to the drift-diffusion model such a change in only the drift rate would result in faster and more accurate judgments. On the other hand, an increase in group size could also result in people being increasingly biased toward the opinion suggested by a group a priori, that is before the actual stimulus is presented. For such a situation, in which only the relative starting point moves, the drift-diffusion model would predict faster judgments for the suggested choice option and slower judgments for the alternative choice option. One can also think of alternative hypotheses, such as a combination of these two mechanisms or an additional inclusion of a third parameter (e.g., the boundary separation). Here, we show that an increase in group size primarily results in people processing the recommended stimuli more efficiently, but also in an increase in response cautiousness. Group size seems to affect this mechanism in a nonlinear way.

Method

Participants

Fifty-one students from the University of Basel, Switzerland, participated in the experiment (*M* = 23.2, *SD* = 3.3 years; 27 female). The study took approximately 2.5 hours after which every participant received 50 CHF plus a variable, performance-contingent bonus of 0.04 CHF for each correct response (*M* = 8.82 CHF; *SD* = 1.32 CHF; range: 7.00-12.00 CHF).

Procedures and Design

Each participant performed a modified version of the face-versus-car categorization task (Philiastides et al., 2011). Participants were informed that they were taking part in a study that aims for a better understanding of the processes underlying object recognition in groups and that the task could be understood as a training session implemented on a popular social media platform. At the beginning of every trial, participants were confronted with cues representing the fictitious opinion of a single individual or a majority of a group of 19 individuals about the category – face or car – of the upcoming stimulus. Next, participants were confronted with dynamic noisy visual stimuli (i.e., short “movies”) of faces or cars and had to decide which of the two was presented. Participants indicated their rating by pressing the appropriate button (counter-balanced left/right between subjects). Because we tried to minimize learning over time, feedback was provided only at the end of the study. We informed participants that cars and faces were presented with equal probability (i.e., 50%). Participants were instructed to solve the task as accurately and as quickly as possible. The experiment started after a short training period.

Pictures of twenty different faces and twenty different cars were used to create dynamic noisy visual stimuli (grayscale images, 8-bit, 256 levels, 500*500 px - see Philiastides et al., 2011 for details) by varying the percentage of phase coherence of each image (see Dakin, Hess, Ledgeway, & Achtman, 2002; Rainer, Augath, Trinath, & Logothetis, 2001 for a detailed description of the algorithms). Each of the 40 stimuli consisted of 30 different frames. The frames were presented at a rate of one frame/50ms and the presentation ended as soon as a participant decided for an option or after 1.5s the latest. Noise was kept constant by using the same percentage of phase coherence across frames.

Social influence was manipulated in three conditions. In two conditions we used visual cues indicating the opinion either of the majority of a large group consisting of 19 fictitious people ("large group condition") or of a single other individual ("small group condition"). In a third condition participants received uninformative visual cues ("control condition" – see Figure 1). To create a more realistic environment in which the opinions of others can also be wrong, opinion cues in the small and large group condition were correct in only 70% of all cases. Importantly, to avoid any biases, opinion cues indicated the correct solution for faces and for cars equally often. The three conditions were presented in a randomized way in three blocks (not to be confused with the three conditions) consisting of twelve mini-blocks with ten stimuli/mini-block. Participants decided in overall 360 trials, in which each mini-block contained stimuli of only one condition (i.e., 12 mini-blocks or 120 stimuli per condition). The three conditions were presented repeatedly in a fixed order of mini blocks. The order was fixed within a participant, but varied between participants. Four different randomizations –

control/large group/small group, large group/small group/control, small group/control/large group and small group/large group/control – were used. After 60, 120, 180 and 300 trials a short self-paced break allowed participants to relax.

Calibration: Noise Level

The noise level was calibrated in a pre-study for every participant individually in a way that aimed for an average accuracy of 60% in the control condition. With an accuracy level of 60% in the control condition, the group opinions with an accuracy of 70% provided useful information that allowed participants to improve their judgments. Thus, participants were incentivized to integrate both sources of evidence, that is, social (opinion cues) and private (dynamic stimuli). Importantly, if participants had blindly followed the opinion cues in the small and/or large group condition, their accuracy would have been 0% for wrong opinions and 100% for correct opinions. Our results (see Accuracies and Response Times) clearly demonstrate that this was not the case.

To calibrate the noise level we used 400 trials consisting of ten noise levels (40 trials/noise level) that were characterized by a different percentage amount of phase coherence (0.1525 - .2650 in steps of .0125, where a lower percentage of phase coherence indicates more noise). After a short training phase, to familiarize participants with the task, different noise levels were presented in mini-blocks of 40 stimuli in a random order with short self-paced breaks in-between. The optimal noise level corresponding to an accuracy level of (approximately) 60% was determined with the *modelfree* R package for fitting psychometric functions (Zychaluk & Foster, 2009).

Method: Drift-diffusion Model Analysis

To examine the role of social influence and group size in rapid perceptual decisions, we estimated the mean drift rate (v), boundary separation (a), starting point (z) and nondecision time (T_{er}) using the drift- diffusion model (Ratcliff, 1978). The drift-diffusion model also allows for between-trial variability in drift (η; normally distributed), starting point (s_z;

uniformly distributed) and nondecision time (s_t; uniformly distributed). We estimated model parameters from the behavioral data from each participant separately, using the Kolmogorov-Smirnov statistic as a fitting criterion (i.e., the maximal vertical distance between observed and predicted cumulative distribution functions of response times), as implemented in the *fast-dm* method (Voss & Voss, 2007). As in Germar et al.'s (2013) previous analyses on social influence, all seven model parameters were allowed to vary freely between the three conditions (control, small group, and large group). The parameters v, a, and z additionally varied as a function of cue type (i.e., the opinion cues, suggesting either face or car) to account for possible perceptual or judgmental bias effects. Finally, separate drift rates were estimated for face and car stimuli (with the lower boundary associated with car decisions and the upper boundary with face decisions), implying 32 parameters per participant (across all conditions and stimuli).

Statistical Data Analysis

Trials in which participants did not respond and trials with RTs < 200ms were excluded from all statistical analyses (4.7% of all trials). We calculated means, standard deviations and $CI_{95\%}$s for the accuracies and response times of all three conditions, separately for wrong and correct opinions: C, $S_{\text{wrong opinions}}$, $S_{\text{correct opinions}}$, $L_{\text{wrong opinions}}$ and $L_{\text{correct opinions}}$.

To test, whether group size (small and large group condition as compared to the control condition) had an effect on accuracy, we calculated the mean, standard deviation, $CI_{95\%}$ and Cohen's d (i.e., μ_{diff}/σ_{diff}) (Cohen, 1988) based on the following effect measure:

$$social\ influence_{gc} = 0.5 \cdot [(A_{gc,co} - A_{cc}) + (A_{cc} - A_{gc,wo})]. \quad (1)$$

Here, A_{cc} refers to accuracy in the control condition and $A_{gc,co}$ and $A_{gc,wo}$ refer to the accuracies in the group conditions (gc) with correct or wrong opinions (co, wo), respectively. The comparison of the response times – wrong versus correct opinions – is based on the following effect measure:

$$influence\ correctness_{oc} = \frac{RT_{sg,oc} + RT_{lg,oc}}{2} - RT_{cc}. \quad (2)$$

RT_{cc} refers to the response time in the control condition and $RT_{sg,oc}$ and $RT_{lg,oc}$ refer to the response times for opinion correctness (oc – correct and wrong opinions) for the small and large group (sg, lg), respectively.

To see, whether $social\ influence_{large\ group}$ > $social\ influence_{small\ group}$ (for the accuracies) and, whether $influence\ correctness_{wrong\ opinion}$ > $influence\ correctness_{correct\ opinion}$ (for the response times), we further calculated the mean, standard deviation, $CI_{95\%}$ and Cohen's d for the difference of these two effect measures (diff. – see Figure 2).

For the diffusion model parameters of interest, that is v, z/a, a (and T_{er}) we calculated the means, standard deviations and $CI_{95\%}$. For the contrasts we again relied on Cohen's d.

Manipulation Checks

At the end of the study we asked participants two questions: (1) "How realistic did you perceive the opinions of the group that we presented to you on the computer screen?" and (2) "How helpful were the opinions of the group that we presented to you on the computer screen?" Answers were provided on a seven-point scale, ranging from 1 (not realistic/not helpful) to 7 (very realistic/very helpful). The results indicate that the opinions were perceived as realistic (M = 4.0, SD = 1.4) and helpful (M = 3.6, SD = 1.4).

Results

We first report the effects of group size and opinion correctness on accuracy and response time to show that both variables effectively influenced participants' behavior. Thereafter, we provide evidence for the appropriateness of the drift-diffusion model by presenting a graphical display of the model's goodness-of-fit. Finally, we focus on the effects of group size and opinion cues on the model parameters. These results illustrate how an increase in group size changes the cognitive processing of categorization judgments.

Accuracies and Response Times

Descriptive measures for accuracy and response time, separately for the different conditions (control, small group and large group) and opinion correctness (no opinion, correct opinion and wrong opinion), are summarized in Table 1.

The accuracy of 59% in the control condition shows that the calibration process worked successfully: Even the average absolute deviation to 60% is according to Cohen's d large, it is ≤10% (*M* = .08, *SD* = .06, $CI_{95\%}$ [.06, .10], *d* = 1.20). Accuracy substantially changed when participants additionally received opinions of groups (see Figure 2A). When receiving correct opinions from a small group the accuracy increased by *M* = .13, *SD* = .12, $CI_{95\%}$ [.09, .16] as compared to the control condition. Correct opinions of large groups (*M* = .18, *SD* = .12, $CI_{95\%}$ [.14, .21]) increased accuracy even more and the comparison between large and small group showed a moderate-large effect of *M* = .05, *SD* = .06, $CI_{95\%}$ [.03, .06], *d* = 0.79. Social influence also reduced the accuracy when a wrong opinion was provided. This negative social influence was again smaller for the small group (*M* = -.15, *SD* = .14, $CI_{95\%}$ [-.19, -.10]) as compared to the large group (*M* = -.20, *SD* = .15, $CI_{95\%}$ [-.24, -.16]) with a moderate effect of *M* = .06, *SD* = .11, $CI_{95\%}$ [.02, .09], *d* = 0.50. These results show that participants used the social information provided by the opinion cues and integrated it with their own private information to make a final judgment. Importantly, participants did not follow the opinion cues blindly as the accuracy is clearly >0% for wrong opinions and <100% for correct opinions. Although the response times for correct opinions did not differ substantially from response times in the control condition (see Figure 2B), response times were larger when participants received wrong opinions as compared to the control condition. Apparently, in a situation with wrong opinions the integration of private and social information can lead to a conflict.

Drift-diffusion Model Fit

For each participant, we estimated 32 parameters with the drift-diffusion model (for the different conditions, stimuli, and cue types – see Method: Drift-diffusion Model Analysis for details). On the basis of these estimated parameter values one can predict the cumulative distribution functions of response times (see Figure 3), separately for all conditions (control, small group, and large group), opinion correctness (wrong versus correct) and stimuli (face and car). Figure 3 illustrates how the models' predictions are related to the observed behavior. This graphical display of model fit (Voss et al., 2013) shows that the drift-diffusion model can qualitatively capture all important aspects of the observed data, that is, the impact of condition, opinion correctness, and stimuli on both accuracy and response times. To quantitatively examine to what extent the model captures the behavioral data, we assessed the models' goodness-of-fit using the product p value from the Kolmogorov-Smirnov tests as fit index (see Voss et al., 2013; Voss, Rothermund, & Voss, 2004). Goodness-of-fit tests with values of $p < .05$ would indicate misfit, which we did not observe for any of the participants (average fit index $M = .42$; $SD = .20$; range: .09 – .88).

Results: Drift-diffusion Model Analysis

In general, the analysis of participants' accuracies and response times clearly demonstrates that the cognitive process underlying the face-versus-car categorization task is affected by group size and opinion correctness. In particular, for accuracies the results show that social influence is stronger in the large as compared to the small group condition. To understand in more detail how these effects can be explained on a psychological level, we modeled the data with the drift-diffusion model. As outlined in the introduction, social influence could specifically affect certain parts of the psychological processing, so we focused our main analyses on three parameters of the model: The drift rate v, the boundary separation a, and the relative starting point a /z. In an explorative step, we also tested whether there is an effect of group size on nondecision T_{er}. The main findings are shown in Figure 3. Table S1 in the Appendix contains a more comprehensive summary of the additional parameters, including the across-trial variabilities.

Drift rate (*v*).

How does an increase in group size affect information uptake? To answer this question, we calculated drift rates for $v_{\text{Control,}}$ $v_{\text{Small Group}}$ and $v_{\text{Large Group}}$ (see Figure 4A left). As the drift rate was allowed to freely vary between stimuli (i.e., was a face or a car presented?) and opinion cues (i.e., was the opinion of the group in favor of face or car?), $v_{\text{Control,}}$ $v_{\text{Small Group}}$ and $v_{\text{Large Group}}$ were derived by calculating mean estimates across stimuli and opinion cues. Very generally, the drift rate can be interpreted as the (relative) amount of evidence accumulated per time unit and increases when a stimuli becomes easier or the perceptual system more sensitive (e.g., Voss et al., 2004). In this experiment, a face (car) judgment is executed when the evidence accumulator hits the upper (lower) boundary and as a consequence drift toward face (car) goes along with a positive (negative) sign. In this particular analysis, we were interested in how an increase in group size affects the efficiency of evidence accumulation independent of which stimuli was presented. Therefore, we averaged the absolute values of the drift rates.

The drift rate in the large group condition (M = 1.01, SD = 0.52, $CI_{95\%}$ [0.87, 1.16]) was higher than in the small group condition (M = 0.83, SD = 0.40, $CI_{95\%}$ [0.72, 0.94]) with a moderate effect between v_{Large} and $v_{\text{Small Group}}$ of M = 0.18, SD = 0.40, $CI_{95\%}$ [0.07, 0.29], d = 0.46. Moreover, there was also a moderate effect between $v_{\text{Small Group}}$ and v_{Control} of M = 0.27, SD = 0.51, $CI_{95\%}$ [0.13, 0.41], d = 0.53, which shows that the drift rate increases from the control condition (M = 0.56, SD = 0.37, $CI_{95\%}$ [0.46, 0.66]) to the small group condition as well. These results show that participants accumulate evidence more efficiently the larger the group. A linear increase can be ruled out, as the increase in group size between $v_{\text{Large Group}}$ and $v_{\text{Small Group}}$ (“+18 opinions”) and $v_{\text{Small Group}}$ and v_{Control} (“+1 opinion”) is not proportional to the increase in drift rate.

If participants' evidence accumulation process was indeed biased towards the choice option favored by the group, we can expect higher positive (negative) drift rates in the large as compared to the small group condition for face (car) opinion cues. To further examine this assumption, we calculated the mean drift rates $v_{\text{Small Face}}$, $v_{\text{Small Car}}$, $v_{\text{Large Face}}$ and $v_{\text{Large Car}}$ for both group sizes and opinion cues by averaging across stimuli (see Figure 4B). Here, in contrast to the previous analysis, we were interested in the direction of the drift, which is why we did not use absolute values. Not surprisingly, the mean drift rates were negative for car cues and positive for face cues. The drift rates for car were more negative (i.e., higher in absolute terms) for the large group condition (M = -0.62, SD = 0.79, $CI_{95\%}$ [-0.84, -0.39]) than for the small group condition (M = -0.41, SD = 0.77, $CI_{95\%}$ [-0.63, -0.19]) with a small effect of M = -0.21, SD = 0.61, $CI_{95\%}$ [-0.38, -0.04], d = 0.34. We found a similar effect for the drift rates for face. Here, $v_{\text{Large Face}}$ (M = 0.89, SD = 0.81, $CI_{95\%}$ [0.66, 1.11]) was more positive than $v_{\text{Small Face}}$ (M = 0.61, SD = 0.68, $CI_{95\%}$ [0.42, 0.81]) with a corresponding small-moderate difference of M = 0.27, SD = 0.57, $CI_{95\%}$ [0.11, 0.43], d = 0.47. The presented analyses clearly show that group size interacts with cue type: Participants' evidence accumulation is biased toward the stimuli cued by the group and this effect is stronger for large as compared to small group opinions.

Finally, we also tested whether one of the two stimuli was processed with more ease by the participants. This was done by comparing v_{Face} with v_{Car} in the control condition. Please note that here the indices "Face" and "Car" correspond to the observed stimuli and not, as in the previous analyses, to the opinions of a group. In the control condition v_{Face} (M = 0.44, SD = 0.55, $CI_{95\%}$ [0.28, 0.59]) was not higher than v_{Car} (M = -0.19, SD = 0.75, $CI_{95\%}$ [-0.40, 0.02]): A within-subjects comparison of the absolute values of v_{Face} and v_{Car} shows that there is no effect (M = -0.01, SD = 0.61, $CI_{95\%}$ [-0.18, 0.16]).

Starting point position (*z/a*).

The relative position of the starting point between the boundaries (i.e., *z* divided by *a*), can be interpreted as a measure of a priori bias toward a decision alternative (e.g., Klauer & Voss, 2008). Values of *z*/*a* larger than 0.5 would indicate a bias toward the choice option associated with the upper boundary (i.e., requiring relatively less information for “upper-boundary” decisions, i.e., faces) and values smaller than 0.5 indicate a bias toward lower-boundary decisions (i.e., cars). The modeling suggests that participants were a priori more in favor of car decisions (with $z/a < 0.5$ in all three conditions – see Figure 4A, middle). The relative starting point decreased with increasing group size and was smaller in the large group condition ($M = 0.41$, $SD = 0.09$, $CI_{95\%}$ [0.39, 0.44]) than in the small group condition ($M = 0.43$, $SD = 0.08$, $CI_{95\%}$ [0.41, 0.45]) and highest in the control condition ($M = 0.45$, $SD = 0.06$, $CI_{95\%}$ [0.43, 0.47]). However, only the contrast $z/a_{\text{Large Group}} - z/a_{\text{Control}}$ shows a small-moderate difference ($M = -0.04$, $SD = 0.08$, $CI_{95\%}$ [-0.06, -0.01], $d = 0.43$), whereas the contrasts $z/a_{\text{Large Group}} - z/a_{\text{Small Group}}$ ($M = -0.02$, $SD = 0.10$, $CI_{95\%}$ [-0.04, 0.01]) and $z/a_{\text{Small Group}} - z/a_{\text{Control}}$ ($M = -0.02$, $SD = 0.07$, $CI_{95\%}$ [-0.04, 0.00]) do not. This analysis shows that the general a priori bias toward cars gets larger with an increase in group size. However, only the difference between the large group and the control condition shows an effect (Figure 4A). Further, there was no interaction between group size and cue type (see Figure 4C).

Boundary separation (*a*).

The drift-diffusion model can map response caution with a larger boundary separation parameter. Increases in *a* lead to slower, but more accurate decisions, as an accidental crossing of the incorrect boundary (i.e., face when car would be correct or vice versa) due to noise in the evidence accumulation process becomes less likely (e.g., Wagenmakers, 2009). Here, an increase in group size resulted in an increase in boundary separation (averaged across cues). Participants were more cautious in their decisions the larger the group; that is, they sacrificed speed for accuracy (see Figure 4A right). The boundary separation in the large group condition (M = 1.28, SD = 0.17, $CI_{95\%}$ [1.24, 1.33]) was higher than the boundary separation in the small group condition (M = 1.23, SD = 0.19, $CI_{95\%}$ [1.18, 1.28]) with a small-moderate effect between $a_{\text{Large Group}}$ and $a_{\text{Small Group}}$ of (M = 0.06, SD = 0.13, $CI_{95\%}$ [0.02, 0.09], d = 0.43). There was also a moderate effect of M = 0.07, SD = 0.13, $CI_{95\%}$ [0.03, 0.10], d = 0.53 between $a_{\text{Small Group}}$ and a_{Control} (M = 1.16, SD = 0.16, $CI_{95\%}$ [1.12, 1.21]). These results provide further evidence for the claim that participants did not just blindly follow the opinions of others (also see Accuracies and Response Times).

Nondecision time (T_{er}).

Although, the nondecision time (T_{er}) parameter is not of main interest in respect to our hypotheses, we nevertheless tested whether an increase in group size leads to a modulation of T_{er}. We found no effect when comparing the large group condition (M = 0.72, SD = 0.21, $CI_{95\%}$ [0.66, 0.77]) with the small group condition (M = 0.72, SD = 0.21, $CI_{95\%}$ [0.66, 0.78]) nor when comparing the small group condition with the control group (M = 0.73, SD = 0.22, $CI_{95\%}$ [0.67, 0.79]): The differences between both $T_{er\ \text{Large Group}}$ and $T_{er\ \text{Small Group}}$ (M = -0.01, SD = 0.06, $CI_{95\%}$ [-0.02, 0.01]) and $T_{er\ \text{Small Group}}$ and $T_{er\ \text{Control}}$ (M = -0.01, SD = 0.06, $CI_{95\%}$ [-0.03, 0.01]) were close to 0. This shows that increases in social influence due to increases in group size do not go along with changes in nondecision processes (e.g., motor preparation).

Discussion

People’s motivation to gain a more accurate perception of reality is a fundamental aspect of social influence and one of the prime reasons why we conform to the opinions of others (Cialdini & Goldstein, 2004; Cialdini & Trost, 1998). Research in the field of social psychology (Asch, 1951, 1955; Bond, 2005; Latané, 1981; Latané & Wolf, 1981; MacCoun, 2012; Mullen, 1983; Tanford & Penrod, 1984) and the *wisdom of the crowds* (Larrick et al., 2012; Mannes, 2009; Surowiecki, 2005) has hotly debated the question of which mathematical model can best describe the functional relationship between group size and (informational) social influence. However, to the best of our knowledge, no previous study has yet applied a sequential sampling model to study the psychological mechanisms that underlie this fundamental relationship. Here, we combined the drift-diffusion model (Ratcliff, 1978) with an adapted version of the face-versus-car categorization task (Philiastides et al., 2011) to compare different hypotheses of how an increase in group size could lead to an increase in social influence. On a behavioral level, we observed that opinion correctness effectively influenced participants' accuracy in the categorization task and that this effect is stronger in the large as compared to the small group condition. Further, we saw that response times increased after participants were confronted with wrong as compared to correct opinions, independent of group size. These results demonstrate that participants were influenced by both the size and the opinion of a group. The drift-diffusion model was able to accurately describe accuracies and response and to transform these measures into psychologically interpretable parameter values. Here, we can see that evidence for options favoured by larger groups was processed more efficiently (effect on *v*), but also that participants required more evidence (effect on *a*) before they made a final decision. Importantly, there was no difference in the relative starting point between the two group conditions. In sum, these results provide additional evidence for the validity of the general social influence mechanism proposed by Germar et al. (2013) and further show, that an increase in group size from one to 19 people modulates this process in a non-linear way.

Of the two bias parameters – drift rate ("perceptual bias") and relative starting point ("judgmental bias") – only the drift rate was modulated by group size and cue type. Further, we also observed an interaction between these two factors, with more positive (negative) drift rates for face (car) cues in the large as compared to the small group condition. Generally, the drift rate is interpreted as the "rate of accumulation of information" (Ratcliff & McKoon, 2008), the "speed of information uptake" (Voss et al., 2013) or the "ease of processing" (Wagenmakers, 2009) and it has been empirically shown that the drift rate decreases when stimuli become harder to discriminate (Ratcliff, 2002; Voss et al., 2004). Drift rate maps task difficulty and/or perceptual sensitivity (e.g., Voss et al., 2004) and when (absolute) drift is high decisions are fast and less driven by noisy fluctuations (Wagenmakers, 2009). Apparently, in this study, evidence for a choice option favoured by a large group was accumulated with more ease than evidence for a choice option favoured by a small group. Importantly, this finding cannot alternatively be explained by a priming effect, as the words “face” or “car” were present on the visual cues of both conditions. It is also noteworthy to mention that effects on the drift rate have also been found in other studies using social information, e.g., in Philiastides and Ratcliff (2013), who studied the effect of branding, or in Klauer et al. (2007), who decomposed the psychological mechanisms of the implicit association test. On a neural level, Philiastides et al. (2011) were able to show that by disrupting the left dorsolateral prefrontal cortex before participants solved the face-versus-car categorization task, the herewith induced decrease in accuracy and increase in response times go along with a decrease in drift rate. This finding points to a neural mechanism, which could also underlie the social influence effect. A potential next step would therefore be to use the adapted version of the face-versus-car categorization task presented here in combination with imaging techniques, such as fMRI.

Besides the effect on the drift rate we also found an increase in boundary separation for large as compared to small groups. The boundary separation parameter is thought to represent “conservatism” or “response cautiousness” and regulates the speed versus accuracy trade-off, where an increase in boundary separation leads to slower, but more accurate decisions (Ratcliff & McKoon, 2008; Voss et al., 2013; Wagenmakers, 2009). It has been shown empirically that accuracy (or speed) instructions can lead to a higher (lower) value of the boundary separation parameter (Ratcliff, 2002; Voss et al., 2004). An increase in boundary separation was also found in the social influence experiments reported by Germar et al. (2013). In our experiment, participants were more cautious in the large as compared to the small group condition and generally more cautious in the group conditions as compared to the control condition, which shows that people do not just blindly follow the opinions of others.

So how do could these two psychological mechanisms work together? In the instructions we explicitly statedthat group opinions can also be wrong and that it is not always good to follow them. Higher response times for wrong as compared to correct opinion cues and accuracies, which were <100% for correct advice and >0% for wrong advice, show that most participants followed this advice and as a consequence integrated both social (opinion cues) and private (noisy stimuli) sources of evidence. The boundary separation is thought to be under subjective control, whereas the drift rate is not (Wagenmakers, 2009). Importantly, when the opinion of a group is inaccurate, a higher drift rate can also lead to faster wrong decisions. As drift rate is most probably not consciously controllable by a person, it could be that participants counteracted the risk of wrong decisions by increasing the amount of evidence necessary to make a final decision. Interestingly, independent of opinion correctness, accuracy in the large group condition (M = .65, SD = .09) is almost equal to accuracy in the small group condition (M = .64 , SD = .09). Mannes (2009) reported that although people generally seem to acknowledge the *wisdom of the crowds* they are not sensitive enough to the information provided by the size of the group. We speculate that this could be due to too strong a scepticism toward large groups as compared to small groups, reflected in a (too) high boundary separation parameter for the large group.

Finally, in our data we can observe that participants responded with “face” slightly more often than with “car” (in 55% of all valid trials). Although we did not find a within-subjects difference when comparing the (absolute) drift rates for face and car in the control condition, the average values show that there could be a tendency to process faces a little bit more easily. Because we mentioned in the instructions that both stimuli will be presented equally often, the a priori bias toward car in all three conditions could reflect a sensible strategy to counteract a tendency to “see” faces more often. An a priori bias can be found in studies where a response leads to greater rewards (Voss et al., 2004) or when two decision options are correct with an unequal probability (Ratcliff & McKoon, 2008). Here, we find a general (a priori) bias toward cars and this bias is higher in the large group as compared to the baseline condition. Why participants increased their bias toward car in the large group condition, however, is not clear to us.

In sum, the results of this study confirm the general findings of Germar et al. (2013) by showing that social influence primarily leads to a more efficient information uptake. Additionally, we can show that this mechanism is scaled by group size. Besides the drift rate, we also found an effect on the boundary separation. These findings extend the knowledge on the psychological mechanism underlying social influence and for the first time show how group size influences the psychological processes of interest.

References

Asch, S. E. (1951). The Effects of Group Pressure Upon the Modification and Distortion of Judgments. In H. Guetzkow (Ed.), *Groups, Leadership and Men: Research in Human Relations* (pp. 177–190). Pittsburgh, PA: Carnegie Press.

Asch, S. E. (1952). *Social Psychology*. Englewood Cliffs, NJ: Prentice-Hall.

Asch, S. E. (1955). Opinions and social pressure. *Scientific American*, *193*, 31–35. doi:10.1038/scientificamerican1155-31

Asch, S. E. (1956). Studies of independence and conformity: A minority of one against a unanimous majority. *Psychological Monographs: General and Applied*. doi:10.1037/h0093718

Aschenbrenner, K. M., Albert, D., & Schmalhofer, F. (1984). Stochastic choice heuristics. *Acta Psychologica*, *56*(1-3), 153–166. doi:10.1016/0001-6918(84)90015-5

Ashby, F. (2000). A Stochastic Version of General Recognition Theory. *Journal of Mathematical Psychology*, *44*(2), 310–329. doi:10.1006/jmps.1998.1249

Baguley, T. (2012). Calculating and graphing within-subject confidence intervals for ANOVA. *Behavior Research Methods, 44*(1), 158–175. doi:10.3758/s13428-011-0123-7

Bond, R. (2005). Group Size and Conformity. *Group Processes & Intergroup Relations*, *8*(4), 331–354. doi:10.1177/1368430205056464

Bond, R., & Smith, P. B. (1996). Culture and conformity: A meta-analysis of studies using Asch's Line judgment task. *Psychological Bulletin*, *119*(1), 111–137. doi:10.1037/0033-2909.119.1.111

Cialdini, R. B., & Goldstein, N. J. (2004). Social influence: compliance and conformity. *Annual Review of Psychology*, *55*, 591–621. doi:10.1146/annurev.psych.55.090902.142015

Cialdini, R. B., & Trost, M. R. (1998). Social influence: Social norms, conformity and compliance. In D. T. Gilbert, S. T. Fiske, & L. Gardner (Eds.), *The Handbook of Social Psychology* (4th ed., pp. 151–192). New York, NY: McGraw-Hill.

Cohen, J. (1988). *Statistical power analysis for the behavioral sciences*. *Statistical Power Analysis for the Behavioral Sciences* (2nd ed., Vol. 2nd, p. 567). Hillsdale, NJ: Lawrence Erlbaum Associates.

Dakin, S. C., Hess, R. F., Ledgeway, T., & Achtman, R. L. (2002). What causes non-monotonic tuning of fMRI response to noisy images? *Current Biology*, *12*(14), R476–R477. doi:10.1016/S0960-9822(02)00960-0

Deutsch, M., & Gerard, H. B. (1955). A study of normative and informational social influences upon individual judgement. *Journal of Abnormal Psychology*, *51*(3), 629–636. doi:10.1037/h0046408

Fehr, E., & Rangel, A. (2011). Neuroeconomic Foundations of Economic Choice — Recent advances. *Journal of Economic Perspectives*, *25*(4), 3–30. doi:10.1257/jep.25.4.3

Germar, M., Schlemmer, A., Krug, K., Voss, A., & Mojzisch, A. (2013). Social Influence and Perceptual Decision Making: A Diffusion Model Analysis. *Personality and Social Psychology Bulletin, 40*(2), 217–231. doi:10.1177/0146167213508985

Gluth, S., Rieskamp, J., & Büchel, C. (2012). Deciding when to decide: time-variant sequential sampling models explain the emergence of value-based decisions in the human brain. *Journal of Neuroscience*, *32*(31), 10686–98. doi:10.1523/JNEUROSCI.0727-12.2012

Gluth, S., Rieskamp, J., & Büchel, C. (2013a). Classic EEG motor potentials track the emergence of value-based decisions. *NeuroImage*, *79*, 394–403. doi:10.1016/j.neuroimage.2013.05.005

Gluth, S., Rieskamp, J., & Büchel, C. (2013b). Deciding not to decide: Computational and neural evidence for hidden behavior in sequential choice. *PLoS Computational Biology, 9*(10). doi:10.1371/journal.pcbi.1003309.

Gold, J. I., & Shadlen, M. N. (2007). The neural basis of decision making. *Annual Review of Neuroscience.*, *30*, 535–574. doi:10.1146/annurev.neuro.29.051605.113038

Guo, F. Y., & Holyoak, K. J. (2002). Understanding similarity in choice behavior: A connectionist model. In *Proceedings of the twenty-fourth annual conference of the cognitive science society* (pp. 393–398). Mahwah, NJ.

Klauer, K. C., & Voss, A. (2008). Effects of Race on Responses and Response Latencies in the Weapon Identification Task: A Test of Six Models. *Personality and Social Psychology Bulletin*, *34*(8), 1124–1140. doi:10.1177/0146167208318603

Klauer, K. C., Voss, A., Schmitz, F., & Teige-Mocigemba, S. (2007). Process components of the Implicit Association Test: a diffusion-model analysis. *Journal of Personality and Social Psychology*, *93*(3), 353–368. doi:10.1037/0022-3514.93.3.353

Laming, D. R. J. (1968). *Information Theory of Choice-Reaction Times.* (p. 172). Oxford, England: Academic Press.

Larrick, R. P., Mannes, A. E., & Soll, J. B. (2012). The social psychology of the wisdom of crowds. In J. I. Krueger (Ed.), *Frontiers of Social Psychology: Social Judgment and Decision Making* (pp. 227–242). New York, NY: Psychology Press.

Latané, B. (1981). The psychology of social impact. *American Psychologist*, *36*(4), 343–356. doi:10.1037/0003-066X.36.4.343

Latané, B., & Wolf, S. (1981). The social impact of majorities and minorities. *Psychological Review*, *88*(5), 438–453. doi:10.1037/0033-295X.88.5.438

Link, S. W., & Heath, R. A. (1975). A sequential theory of psychological discrimination. *Psychometrika*, *40*(1), 77–105. doi:10.1007/BF02291481

MacCoun, R. J. (2012). The burden of social proof: Shared thresholds and social influence. *Psychological Review*, *119*(2), 345–372. doi:10.1037/a0027121

Mannes, A. E. (2009). Are We Wise About the Wisdom of Crowds? The Use of Group Judgments in Belief Revision. *Management Science*, *55*(8), 1267–1279. doi:10.1287/mnsc.1090.1031

Mullen, B. (1983). Operationalizing the effect of the group on the individual: A self-attention perspective. *Journal of Experimental Social Psychology*, *19*(4), 295–322. doi:10.1016/0022-1031(83)90025-2

Nosofsky, R. M., & Palmeri, T. J. (1997). An exemplar-based random walk model of speeded classification. *Psychological Review, 104*(2), 266–300. doi:10.1037/0033-295X.115.2.446

Philiastides, M. G., Auksztulewicz, R., Heekeren, H. R., & Blankenburg, F. (2011). Causal Role of Dorsolateral Prefrontal Cortex in Human Perceptual Decision Making. *Current Biology*, *21*(11), 980–983. doi:10.1016/j.cub.2011.04.034

Philiastides, M. G., & Ratcliff, R. (2013). Influence of branding on preference-based decision making. *Psychological Science*, *24*(7), 1208–15. doi:10.1177/0956797612470701

Rainer, G., Augath, M., Trinath, T., & Logothetis, N. K. (2001). Nonmonotonic noise tuning of BOLD fMRI signal to natural images in the visual cortex of the anesthetized monkey. *Current Biology*, *11*(11), 846–854. doi:10.1016/S0960-9822(01)00242-1

Ratcliff, R. (1978). A theory of memory retrieval. *Psychological Review*, *85*(2), 59–108. doi:10.1037/0033-295X.85.2.59

Ratcliff, R. (2002). A diffusion model account of response time and accuracy in a brightness discrimination task: Fitting real data and failing to fit fake but plausible data. *Psychonomic Bulletin & Review*, *9*(2), 278–291. doi:10.3758/BF03196283

Ratcliff, R., & McKoon, G. (2008). The Diffusion Decision Model: Theory and Data for Two-Choice Decision Tasks. *Neural Computation*, *20*(4), 873–922. doi:10.1162/neco.2008.12-06-420

Rieskamp, J. (2008). The probabilistic nature of preferential choice. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *34*(6), 1446–1465. doi:10.1037/a0013646

Roe, R. M., Busemeyer, J. R., & Townsend, J. T. (2001). Multialternative decision field theory: A dynamic connectionst model of decision making. *Psychological Review*, *108*(2), 370–392. doi:10.1037/0033-295X.108.2.370

Smith, P. L. (1995). Psychophysically principled models of visual simple reaction time. *Psychological Review*, *102*(3), 567–593. doi:10.1037/0033-295X.102.3.567

Surowiecki, J. (2005). *The Wisdom of Crowds*. New York, NY: Anchor Books.

Tanford, S., & Penrod, S. (1984). Social Influence Model: A formal integration of research on majority and minority influence processes. *Psychological Bulletin*, *95*(2), 189–225. doi:10.1037/0033-2909.95.2.189

Usher, M., & McClelland, J. L. (2001). The time course of perceptual choice: The leaky, competing accumulator model. *Psychological Review*, *108*(3), 550–592. doi:10.1037/0033-295X.108.3.550

Usher, M., & McClelland, J. L. (2004). Loss aversion and inhibition in dynamical models of multialternative choice. *Psychological Review*, *111*(3), 757–769. doi:10.1037/0033-295X.111.3.757

Vickers, D. (1979). *Decision Processes in Visual Perception* (p. 406). New York, NY: Academic Press.

Voss, A., Nagler, M., & Lerche, V. (2013). Diffusion models in experimental psychology: A practical introduction. *Experimental Psychology*, *60*, 385–402. doi:10.1027/1618-3169/a000218

Voss, A., Rothermund, K., & Voss, J. (2004). Interpreting the parameters of the diffusion model: an empirical validation. *Memory & Cognition*, *32*(7), 1206–1220. doi:10.3758/BF03196893

Voss, A., & Voss, J. (2007). Fast-dm: a free program for efficient diffusion model analysis. *Behavior Research Methods*, *39*(4), 767–775. doi:10.3758/BF03192967

Wagenmakers, E.-J. (2009). Methodological and empirical developments for the Ratcliff diffusion model of response times and accuracy. *European Journal of Cognitive Psychology*, *21*(5), 641–671. doi:10.1080/09541440802205067

Wallsten, T. S., & Barton, C. (1982). Processing probabilistic multidimensional information for decisions. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *8*(5), 361–384. doi:10.1037/0278-7393.8.5.361

Zychaluk, K., & Foster, D. H. (2009). Model-free estimation of the psychometric function. *Attention, Perception & Psychophysics*, *71*(6), 1414–1425. doi:10.3758/APP.71.6.1414

Table 1

Accuracy and Response Time

		Accuracy		Response Time	
condition	opinion correctness	*M* (*SD*)	$CI_{95\%}$	*M* (*SD*)	$CI_{95\%}$
control	no opinions	.59 (.10)	[.56, .62]	977 (204)	[920, 1035]
small	correct	.72 (.12)	[.68, .75]	979 (207)	[920, 1037]
	wrong	.44 (.17)	[.39, .49]	998 (205)	[940, 1055]
large	correct	.76 (.13)	[.73, .80]	974 (204)	[917, 1031]
	wrong	.39 (.18)	[.34, .44]	995 (215)	[935, 1056]

Notes. Accuracies represent the group mean of the mean accuracies of each participant. Response times represent the group mean of the median response times of each participant. CI = confidence interval. For all within-subjects conditions we had a sample size of $N = 51$.

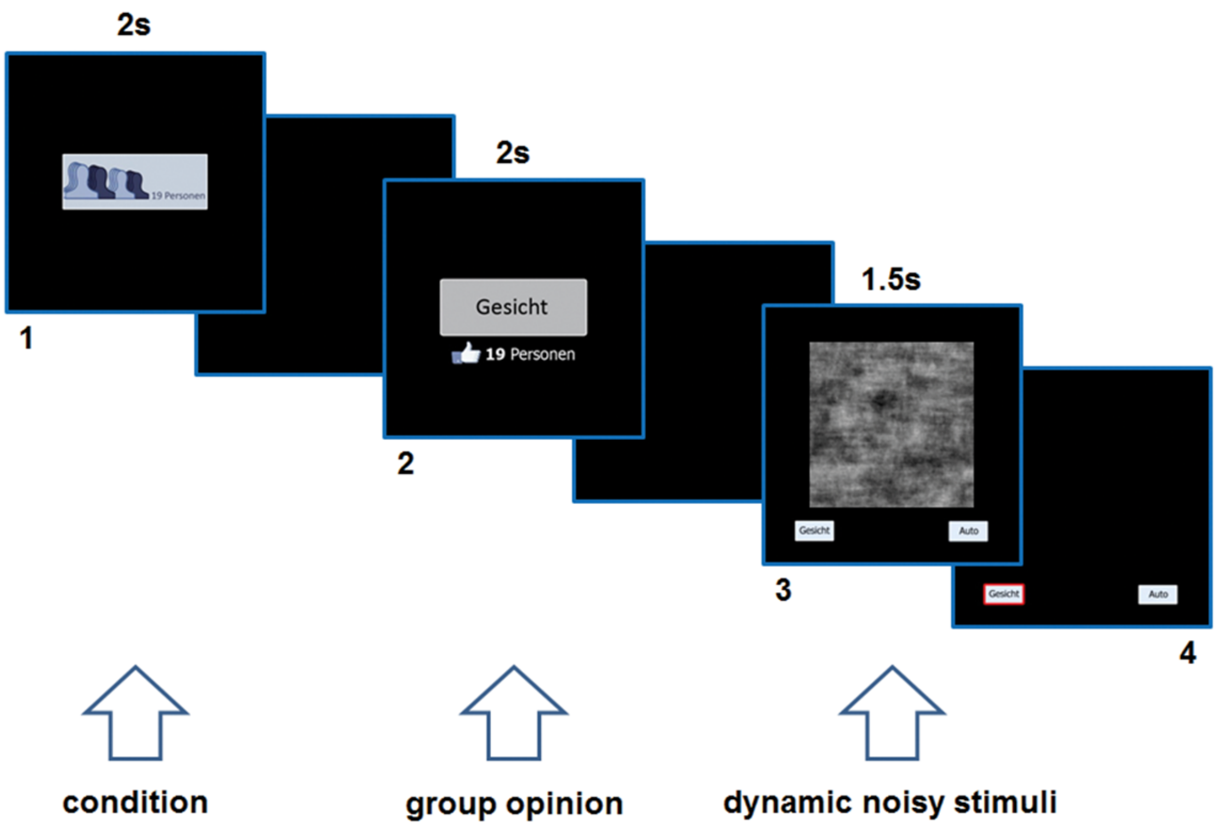

Figure 1. Trial structure. (1) At the beginning of each mini-block the condition (control, small group or large group) was indicated. (2) At the beginning of every trial, participants were exposed to a visual cue indicating the opinion (“face” or “car”) of the majority of a group consisting of 19 individuals (large group), a visual cue indicating the opinion of a single individual (small group) or a visually analogous cue with no further information (control). For the visual cue of the control condition we replaced the group opinion – “face” (“Gesicht” in German) or “car” (“Auto”) – with “Xxxx” and the number of persons below the opinion was set to 00 instead of 01 or 19. (3) Dynamic noisy visual stimuli of either faces or cars were presented for a maximum of 1.5s during which participants made a decision.

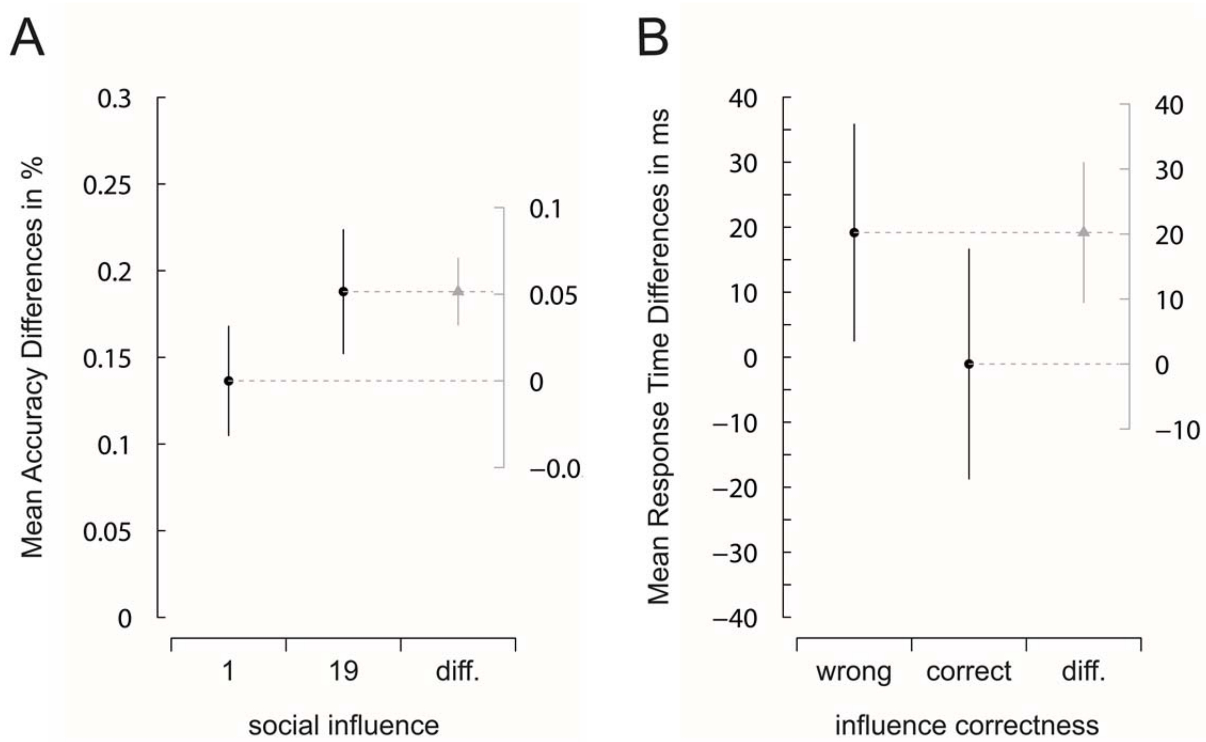

Figure 2. Effects of group size and opinion correctness on accuracy and response time: (A) $social\ influence_{gc}$ (see methods for details) for accuracy was calculated for the small (1) group (*M* = .14, *SD* = .11, $CI_{95\%}$ [.10, .17], *d* = 1.22) and for the large (19) group condition (*M* = .19, *SD* = .13, $CI_{95\%}$ [.15, .22], *d* = 1.48). The contrast $social\ influence_{large\ group}$ - $social\ influence_{small\ group}$ (diff.) shows a moderate-large main effect between groups on accuracy of *M* = .05, *SD* = .07, $CI_{95\%}$ [.03, .07], *d* = 0.75. (B) The response time $influence\ correctness_{wrong\ opinion}$ (*M* = 19ms, *SD* = 59ms, $CI_{95\%}$ [3ms, 36ms], *d* = 0.32) was larger than the $influence\ correctness_{correct\ opinion}$ (*M* = -1ms, $95_{\%}$, *SD* = 63ms, $CI_{95\%}$ [-19ms, 17ms]). This corresponds to a moderate main effect of opinion correctness on response times (diff.) of *M* = 20ms, SD = 38ms, $CI_{95\%}$ [9ms, 31ms], *d* = .53.

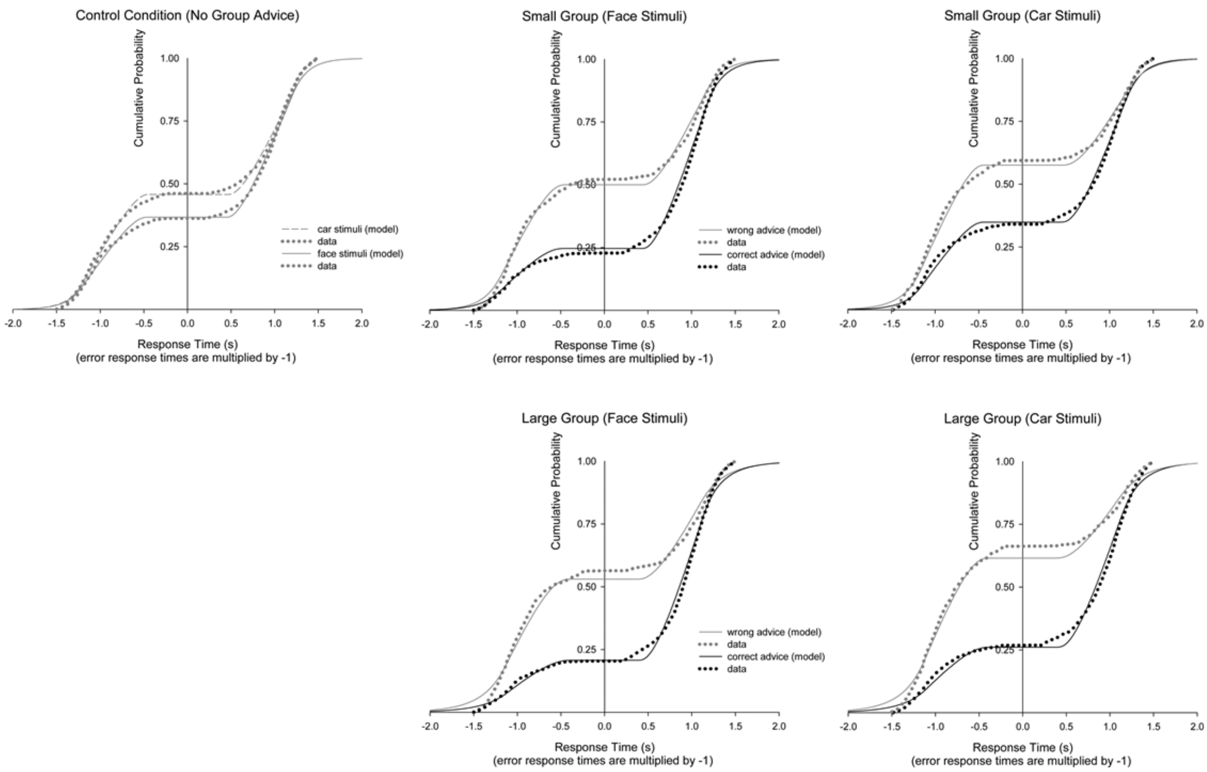

Figure 3. Graphical displays of model fit (based on aggregated data, averaged across participants). The plots show the observed (dotted lines) and predicted (straight lines) cumulative distribution functions of response times, as function of condition (control, small group, and large group), opinion correctness (correct and wrong opinions), and observed stimulus (face or car). Error response times were multiplied by –1 and are displayed on the negative side of the x-axis. The point where a curve intersects with the y-axis represents the proportion of error responses.

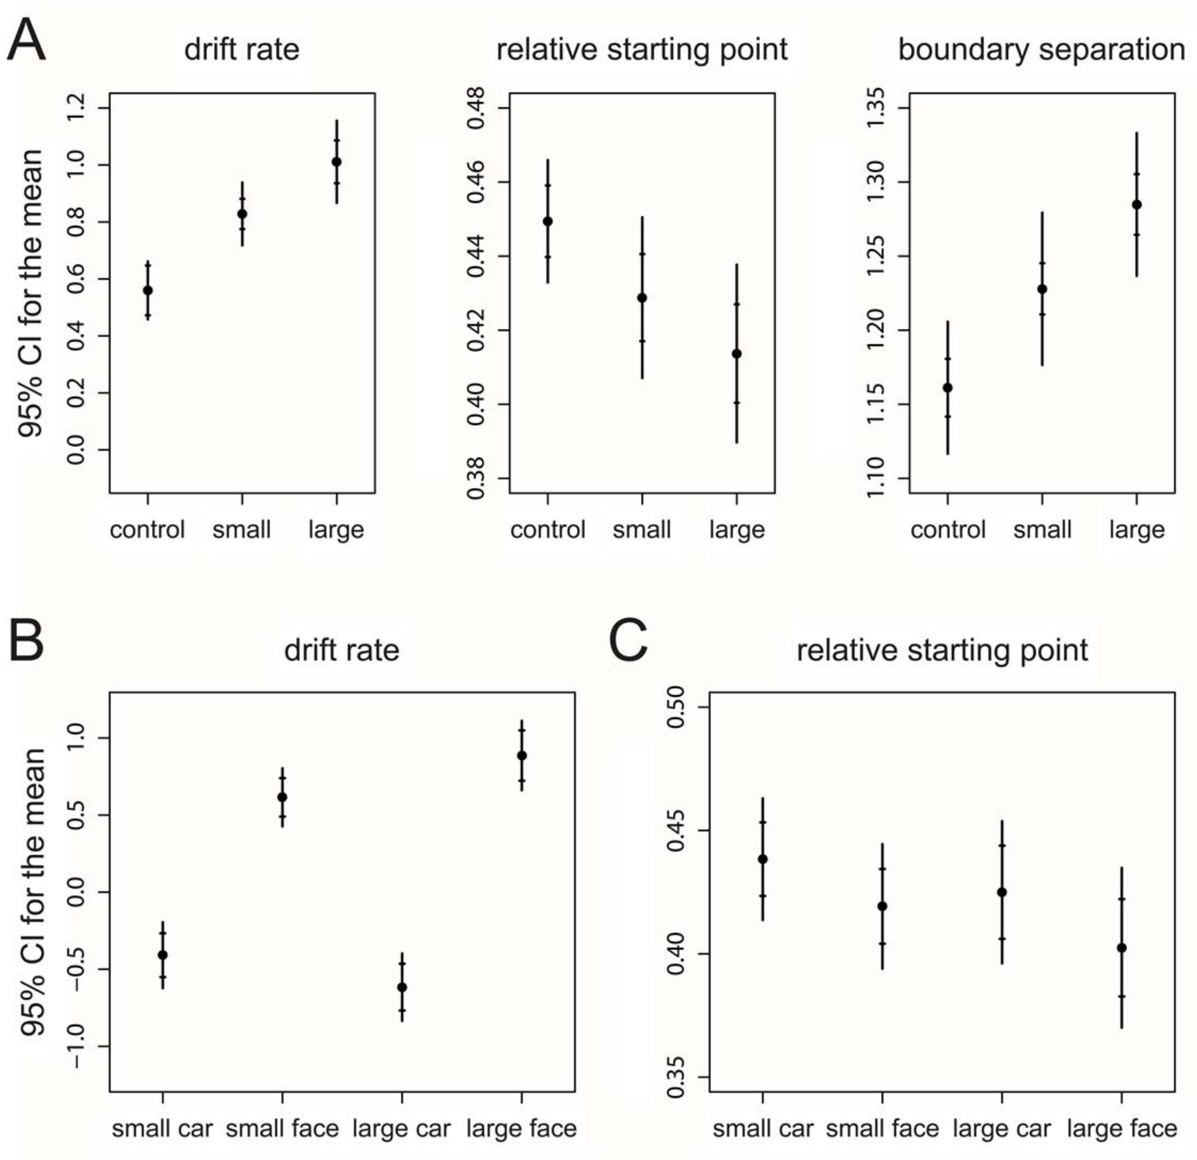

Figure 4. The effect of group size and opinion cues on information processing. Displayed are the width-adjusted Cousineau-Morey $CI_{95\%}$s (inner-tiers) and the multilevel $CI_{95\%}$s (outer-tiers) for (A) the drift rate (left), the relative starting point (middle) and the boundary separation (right) separately for the control, the small group and the large group condition (B) the drift rate and (C) the relative starting point depending on opinion cue ("car" and "face") and group size (small and large). Width-adjusted Cousineau-Morey (CM) and Multilevel (ML) $CI_{95\%}$s were computed using the R functions provided by Baguely (2012). In a within-subjects design non-overlap of the CM $CI_{95\%}$s (inner tiers) for two means corresponds to a $CI_{95\%}$ of the difference between the two means, which doesn't include 0. ML $CI_{95\%}$s (outer tiers) are of interest, when one wants to see, if a particular parameter value is a plausible

candidate for the estimated mean (e.g. is a relative starting point of .5 a plausible value in condition XY).

Appendix

Table S1

Additional parameters of the diffusion model

parameter	*M* (*SD*)	$CI_{95\%}$
$z_{Control}$	0.52 (0.06)	[0.50, 0.53]
$z_{Small\ Car}$	0.52 (0.09)	[0.49, 0.54]
$z_{Small\ Face}$	0.51 (0.09)	[0.49, 0.54]
$z_{Large\ Car}$	0.54 (0.16)	[0.49, 0.59]
$z_{Large\ Face}$	0.51 (0.16)	[0.46, 0.55]
$a_{Control}$	1.16 (0.16)	[1.12, 1.21]
a_{Small}	1.23 (0.19)	[1.18, 1.28]
a_{Large}	1.28 (0.17)	[1.24, 1.33]
$a_{Small\ Car}$	1.20 (0.20)	[1.15, 1.26]
$a_{Small\ Face}$	1.25 (0.19)	[1.20, 1.30]
aLarge Car	1.29 (0.21)	[1.23, 1.34]
aLarge Face	1.28 (0.20)	[1.23, 1.34]
$v_{Small\ Car\ Car}$	-0.70 (0.93)	[-0.96, -0.44]
$v_{Small\ Car\ Face}$	0.25 (0.93)	[-0.01, 0.51]
$v_{Small\ Face\ Car}$	-0.12 (0.87)	[-0.36, 0.13]
$v_{Small\ Face\ Face}$	0.98 (0.66)	[0.80, 1.17]
$v_{Large\ Car\ Car}$	-0.99 (0.81)	[-1.22, -0.77]
$v_{Large\ Car\ Face}$	0.59 (0.97)	[0.31, 0.86]
$v_{Large\ Face\ Car}$	-0.24 (0.97)	[-0.51, 0.03]
$v_{Large\ Face\ Face}$	1.19 (0.91)	[0.93, 1.44]
$\eta_{Control}$	0.48 (0.27)	[0.41, 0.56]
η_{Small}	0.52 (0.28)	[0.44, 0.59]
η_{Large}	0.58 (0.25)	[0.51, 0.65]
$s_{zControl}$	0.33 (0.12)	[0.30, 0.36]
s_{zSmall}	0.34 (0.10)	[0.31, 0.37]
s_{zLarge}	0.36 (0.08)	[0.34, 0.39]
$s_{tControl}$	0.02 (0.01)	[0.01, 0.02]
s_{tSmall}	0.02 (0.01)	[0.02, 0.02]
s_{tLarge}	0.02 (0.01)	[0.01, 0.02]

Notes. All values are based on a sample size of $N = 51$. The first index refers to group size (Control, Small and Large). For all parameters – except *v* – Car or Face refers to the opinion cue. For *v* the second index refers to the observed stimuli (Car versus Face), whereas the third index refers to the opinion cue (Car versus Face).

Title: Neural Correlates of Informational Cascades: Brain Mechanisms of Social Influence on Belief Updating

Running title: Neural Correlates of Informational Cascades

Authors: Rafael E. Huber[1], Vasily Klucharev[1,2, 3], and Jörg Rieskamp[1]

Affiliation:

[1] Department of Psychology, University of Basel, Missionsstrasse 62A, 4055 Basel, Switzerland

[2] Department of Higher Nervous Activity and Psychophysiology, Saint Petersburg State University199034, University Emb.7-9, Saint Petersburg, Russia

[3] Department of Psychology, National Research University Higher School of Economics, 46B Volgogradskiy Prospekt, Moscow, Russia

Corresponding author: Rafael E. Huber, University of Basel, Department of Psychology, Missionsstrasse 62A, 4055 Basel, Switzerland, Phone: +41 (0)61 267 02 98, Fax: +41 (0)61 267 06 13, rafael.huber@unibas.ch

Number of pages: 36

Number of figures: 5

Number of tables: 2

Number of words for Abstract: 195

Number of words for Text: 5000

Funding: This work was supported by the Swiss National Science Foundation (SNSF grant no. 100014-130352) to Vasily Klucharev and Jörg Rieskamp and by the Switzerland–Russian Scientific & Technological Cooperation Program, the Russian Targeted Federal Program "Scientific and scientific-pedagogical personnel of innovative Russia" (contract 8488, by a grant from the Russian Foundation for Basic Research [RFBR 11-06-00449-a]).

Acknowledgments: We thank Markus Klarhoefer, Oliver Schürmann, and Stefan Thommen for assistance in the fMRI experiments and Andreas Pedroni and Peter Mohr for helpful comments.

Conflict of Interest: The authors declare no competing financial interests.

Please note: A final version of this manuscript has been accepted for publication in Social Cognitive and Affective Neuroscience:

Huber, R. E., Klucharev, V., & Rieskamp, J. (2014). Neural Correlates of Informational Cascades: Brain Mechanisms of Social Influence on Belief Updating. *Social Cognitive and Affective Neuroscience* . doi:10.1093/scan/nsu090

Abstract

Informational cascades can occur when rationally acting individuals decide independently of their private information and follow the decisions of preceding decision makers. In the process of updating beliefs, differences in the weighting of private and publicly available social information may modulate the probability that a cascade starts in a decisive way. By using functional magnetic resonance imaging, we examined neural activity while participants updated their beliefs based on the decisions of two fictitious stock market traders and their own private information, which led to a final decision of buying one out of two stocks. Computational modeling of the behavioral data showed that a majority of participants overweighted private information. Overweighting was negatively correlated with the probability of starting an informational cascade in trials especially prone to conformity. Belief updating by private information was related to activity in the inferior frontal gyrus/anterior insula, the DLPFC, and the parietal cortex; the more a participant overweighted private information, the higher the activity in the inferior frontal gyrus/anterior insula and the lower in the parietal-temporal cortex. This is the first study exploring the neural correlates of overweighting of private information, which underlies the tendency to start an informational cascade.

INTRODUCTION

Research in the social sciences has reliably demonstrated that individuals are influenced by the behavior of others (e.g., Cialdini and Goldstein, 2004; Raafat, Chater, and Frith, 2009). Stock market bubbles, for example, can emerge when traders start to follow misleading decisions made by their colleagues, disregarding their own private information. Interestingly, theoretical and empirical work in economics has shown that initial decisions of others can create an environment in which it is even rational for subsequent decision makers to disregard their own private information and to follow others. Such a pattern of conforming decisions is called an informational cascade (Anderson and Holt, 1997; Banerjee, 1992; Bikhchandani *et al*., 1992). Usually, informational cascades lead to a desired outcome. However, a "reverse" cascade can arise if a substantial number of initial decision makers receive an incorrect private signal and therefore make incorrect decisions. In such situations, all subsequent decision makers would rationally follow the initial decisions and ignore their own private signals. The theory of informational cascades can explain numerous real-life phenomenon, such as nonemployment in the labor market (Oberholzer-Gee, 2008), revolutionary regime transitions (Ellis and Fender, 2011), and financial crises (Chari and Kehoe, 2004). The probability that a cascade starts strongly depends on how people weight and integrate their own private as compared to publicly available social information (Bernardo and Welch, 2001; Goeree *et al*., 2007; Nöth and Weber, 2003). Weizsäcker's (2010) meta-analysis suggests that people tend to overweight private as compared to social information, even in situations in which following others is beneficial. Due to overweighting of private information, cascades might occur less often, as predicted by the theory of informational cascades. In the present work, we combine neurobiological, economic, and computational approaches to investigate the neural mechanism of (biased) belief updating during financial decisions and to explore individual differences in the weighting and processing of private information, which can modulate the frequency of starting a cascade.

From a cognitive perspective, informational cascades are based on a process of sequential belief updating of social and private information, on which a final decision under uncertainty rests. Recent studies in the field of decision neuroscience provide evidence for the involvement of the anterior insula (Preuschoff *et al*., 2006, 2008), the anterior insula in combination with the inferior frontal gyrus (Paulus *et al*., 2003), the posterior fronto-median cortex (Volz *et al*., 2003, 2004) and the parietal cortex, often in combination with the dorsolateral prefrontal cortex (Huettel et al., 2005; Mohr *et al*., 2010; Stern *et al*., 2010; Symmonds *et al*., 2011; Vickery and Jiang, 2009; Wright *et al*., 2012), in belief updating and decision making under uncertainty (see (Bach *et al*., 2011) for an overview). Whereas the inferior parietal lobule (angular gyrus) seems to have a special role in tracking observed relative frequencies of events, activity within a region of the inferior frontal gyrus has been found to be negatively correlated with Bayesian posterior probability (d'Acremont *et al.*, 2013).

Contrary to other paradigms exploring belief updating (e.g., the evidence accumulation task; Stern *et al*., 2010 or the ball/bin betting task by d'Acremont *et al.*, 2013), informational cascades require people not only to update a belief on the basis of (private) information, but additionally to derive social information from the observed decisions of others. A better understanding of the differences in updating private as compared to social information is crucial for the theory of informational cascades, because overweighting of private information can result in fewer cascades than predicted by the theory. Here, for the first time, we investigate the neural mechanism of biased belief updating of private as compared to social information.

MATERIALS AND METHODS

Participants

Thirty-two people recruited from the subject pool of the University of Basel participated in our experiment. Five participants were excluded from the final data analysis (two because of technical problems during the fMRI data acquisition, one because of a technical error in the experimental script, one because of misuse of the response device, and one because of left-handedness). The final sample consisted of 27 healthy, right-handed participants with normal or corrected-to-normal vision (mean age = 22.4 years, ± 2.0 years SD, 20-29 years, 9 females). The study was approved by the local ethics committee and participants gave written informed consent. Participation in the study was reimbursed with a fixed amount of 30 CHF and a variable bonus (mean bonus = 3.99 CHF, ± 0.42 CHF SD, 2.90-4.60 CHF). The variable bonus was performance contingent, so that deviations from the correct probability estimate led to a lower bonus following a non-linear quadratic scoring rule (Selten, 1998).

Experimental design

We used a hypothetical decision scenario representing an adapted version of the classical informational cascades paradigm (Anderson and Holt, 1997). In our study, participants acting as stock market traders were required to repeatedly choose the profitable ("$good$") of two stocks (W or S) given some evidence e. Participants were told that stock markets are very volatile and fast moving and that every week (trial) only one stock is profitable. At the end of each trial, participants reported the posterior probability $p(good|e)_t$ that the chosen stock was "$good$" (Figure 1). In the 32 experimental trials, participants sequentially received three different pieces of evidence. At the beginning of a trial, two decisions made by other fictitious traders (trader I and II) in the "Swiss Capital Bank" were shown, representing *social information I* and *social information II*. The *social information* was followed by *private information* in the form of a personal recommendation from a rating agency. Participants were informed that all other traders also received their own personal recommendation from an independent rating agency. The likelihood $p(e|good)_t$ of receiving a correct recommendation from a rating agency was 2/3 (indicated by the visual cue: “+”) or 4/5 (visual cue: “++”) for all traders and for the participant. The quality (“+” or “++”) of the recommendations received was indicated on the screen above the decisions of the other traders (*social information I* and *II*) or above the *private information* for the participant. The posterior probability that one of the two stocks was profitable (“$good$”) given the received and perceived evidence can be determined following Bayes theorem as:

$$p(good|e)_t = \frac{p(good|e)_{t-1} \cdot p(e|good)_t}{p(good|e)_{t-1} \cdot p(e|good)_t + p(bad|e)_{t-1} \cdot p(e|bad)_t}, \tag{1}$$

where *t* refers to the three different points in time in the belief updating process (see Figure 1). At *t* = 0 without a participant having received any information $p(good|e)_{t=0}$ = 0.50. Based on the assumption that other traders incorporated all available evidence, participants could derive the recommendation received by other traders. Because trader I always received low (“+”) quality recommendations her decision (*social information I*) signaled the correct stock with a likelihood of 0.67 (i.e., $p(e|good)$ = 0.67). Next, trader II was confronted with a recommendation of either low (“+”) quality (i.e., $p(e|good)$ = .67) or high (“++”) quality (i.e., $p(e|good)$ = .80). This evidence could then be combined with the information inferred from the decision of the first trader, which led to four possible posterior probabilities of the chosen stock by trader II (i.e., 0.50; 0.67; 0.80; 0.89). After receiving a personal recommendation (*private information*) participants could update their belief, which should correspond to six different posterior probabilities (i.e., 0.50; 0.67; 0.80; 0.89; 0.94; and 0.97). Importantly, by using all different combinations of decisions and private information (2×4×4 = 32 trials of interest), we created a design matrix in which the different pieces of evidence are independent; that is, seeing one piece of evidence did not allow the prediction of the next piece of evidence. To force participants to pay equal attention to *social* and *private information* and to update their probability estimate at every point in time (t), we included six filler trials in the task. In these trials, subjects had to make a decision with only one (*social information I*) or two (*social information I & II*) pieces of evidence and no *private information*. To familiarize themselves with the task, participants completed 11 training trials outside of the scanner before the fMRI session. To further boost their attention, filler trials were overrepresented in these training trials. The randomized sequence of trials was identical for all subjects. Trials were separated with fixation crosses, as were the different events within a trial (see Figure 1). The inter-stimulus intervals (ISI) between the time windows were varied according to a left truncated Poisson distribution (mean (λ) = 3172.78 ms, min = 1000 ms, max = 8000 ms). Importantly, from a normative Bayesian perspective, the first two decision makers can create a situation in which the third decision maker (and all subsequent decision makers) should ignore *private information* and just follow the decisions of others. Thus, the decision of the third decision maker is crucial, as it can start or prematurely end an informational cascade. Therefore, in our paradigm we investigated the cognitive and neural mechanisms underlying the process of belief updating and decision making of the third decision maker, who can initiate or end an informational cascade.

Behavioral data analysis

To examine whether participants differentiated between the six different posterior probabilities (i.e., $p(good|e)_{t=3}$ = 0.50, 0.67, 0.80, 0.89, 0.94, and 0.97), we performed a one-way repeated measures ANOVA with the six levels of uncertainty as within-subject factor and the average probability judgments as the dependent variable. The same analysis was conducted with the logarithm of the reaction times as dependent measure.

Conformity index

The experimental design matrix included six "conflicting" trials in which the two pieces of *social information* suggested buying the same stock whereas the *private information* suggested buying the other stock and where the normatively correct decision was consistent with the *social information* and opposite to the *private information*. Therefore, we calculated a *conformity index* for every participant, defined as the percentage of decisions in line with the decision of the others in these specific trials.

Computational models

To explain the cognitive process underlying belief updating, we constructed an *Evidence Model* that represents a modification of the model proposed by Hung and Plott (2001). According to the normative Bayesian solution (see Equation 1), a participant is required to update her prior belief with every new piece of evidence e_t presented at *t*. To simplify the Bayesian solution, Equation 1 can be transformed by computing the log odds ratio of the posterior probabilities of which of the two stocks being the profitable one ("$good$") assuming equal priors (e.g., Dieckmann and Rieskamp, 2007); that is,

$$ln \frac{p(good_W|e_t)}{p(good_S|e_t)} = \sum_{t=1}^{T} ln \frac{p(e_t|good_W)}{p(e_t|good_S)} \tag{2}$$

However, people might not follow the Bayesian solution and might weight their *private information* more heavily than the socially inferred information. To identify how people weight the different pieces of information, we extended Equation 2 by allowing pieces of information to be weighted differently; that is,

$$\hat{Y} = \beta_0 + \sum_{t=1}^{T} \beta_t \cdot ln \frac{p(e_t|good_W)}{p(e_t|good_S)} \tag{3}$$

where β_0 represents a bias for one of the two stocks at $t = 0$ and β_t refers to the weight given to the different pieces of information. If all weights are equal to 1 and $\beta_0 = 0$ then Equation 3 is identical to Equation 2; that is, the normative solution is nested within the Evidence Model specified by Equation 3.

When estimating the Evidence Model (see supplementary methods), we also imposed three different constraints on the model parameters. First, in the full model (FM) we estimated one bias parameter β_0 and three different β_t weights for each piece of information at the three points in time (*social information I, social information II,* and *private information*), providing four parameters. Second, for the social model (SM), we assumed no bias (i.e., $\beta_0 = 0$) and one single weight for *social information* (i.e., $\beta_{t=1} = \beta_{t=2}$) and one weight for *private information* (i.e., $\beta_{t=3}$), leading to a total of two free parameters. Third, we also determined the goodness-of-fit of the normative Bayesian model (BM) by setting $\beta_0 = 0$ and all other weights to one (i.e., $\beta_t = 1$). Whereas the BM has no flexibility in weighting information differently, the FM allows weighting each piece of information in a different way. The SM assumes that people do not have a bias for one of the options, treat both pieces of *social information* equally but weight their *private information* differently. The SM is more complex than the BM but less complex than the FM.

Information weighting index

A decision maker following Bayesian principles should weight the *social* and *private information* equally. To examine to what extent participants deviated from the Bayesian approach, we determined an *information weighting index* for the SM by dividing the estimated weight for the *private information* ($\beta_{t=3}$, i.e., using the mode of the marginal posterior distribution as a point estimate) by the sum of the estimated weights for the *private* and *social information* (i.e., $\beta_{t=3} + \beta_{t=1+2}$). An *information weighting index* > 0.50 indicates overweighting of *private* as compared to *social information*, whereas values of < 0.50 indicate overweighting of *social* as compared to *private information*.

Functional imaging data analyses

To study the neural underpinnings of belief updating with social and private information, two first level models were calculated in the context of a GLM (SPM8, Wellcome Trust Center for Neuroimaging, University College London). Our experimental design is characterized by three updating stages (see Figure 1). In every trial, participants were forced to update their belief $p(good|e)_t$ after the decisions of two traders (*social information I & II)* and after they had received their own *private information*.

We computed how much a signal given at $t = 2$ increased/decreased the belief in the option that was more probable at stage $t = 1$ following the Bayesian solution (i.e., Equation 1). Likewise, we determined the difference of the posterior probability between $t = 2$ and $t = 3$. Please note that as the decision of trader 1 was always based on a low (+) quality signal for either stock *W* or *S*. Belief updating from $t = 0$ (i.e., the beginning of a trial) to $t = 1$ was the same for every trial and therefore not explicitly modeled.

First level analysis

In the *first level model 1*, belief updating at the *social information II* (belief updating by *social information*) and at the *private information* (belief updating by *private information*) stages was modeled with a single parametric regressor to account for general effects of belief updating at both stages (i.e., independent of the social or private nature of the information). Brain activity at the time of the decision and at the time of the probability judgment was modeled with separate parametric regressors tracking the log odds of the probability judgments and the decision for either stock *W* or *S*. We also included parametric regressors coding for the stock with the highest posterior probability (at $t = 1$ and $t = 2$ and 3 combined) and for the quality of the *private information* (low (+) or high (++)) at $t = 2$ and 3 combined. Decision and/or probability judgment time windows in which participants gave no answer and filler stimuli were included in the GLM as regressors of no interest.

In the *first level model 2,* the second (*social information II*) and third (*private information*) belief updating stages were modeled separately using parametric regressors to account for the specific effects of belief updating by *social* and *private information*. The quality of the *private information* (low (+) or high (++)) was included as a parametric regressor for the belief updating stage at $t = 3$. In all other respects, first level models 1 and 2 were similar. To account for head movements, both first level models included motion parameters.

Second level analysis

To test for the general (*first level model 1*) and specific (*first level model 2*) effects of belief updating as well as for the effects of an increase in subjective uncertainty during decision making (*first level model 2* – see supplementary fMRI results) we used one-sample *t*-tests on the group level ($P < 0.001$ (uncorrected) with a minimum cluster size of 20 voxels). To test how belief updating by private information was modulated by inter-individual differences in information weighting, we used a multiple regression design ($P < 0.001$ or 0.005; uncorrected) with the *information weighting index* as a covariate. In order to restrict the search volume only to brain regions involved in belief updating by *private information* we used the results of the respective second level analysis as an explicit mask (using a liberal threshold, $P < 0.005$, uncorrected, for the mask). To further illustrate these findings we extracted the contrast estimates within two ROIs (see Figure 5) and plotted them against the *information weighting index*. The ROIs were defined with the MarsBaR toolbox for SPM (Brett *et al.*, 2002).

RESULTS

Behavioral results

Overall, participants performed the task consistent with the Bayesian solution: In 93.18% of all trials in which participants (N = 27) made a decision, they decided in accordance with the Bayesian solution, with seven participants always choosing the more profitable stock. The six different levels of uncertainty significantly modulated participants' probability judgments, F(3.64, 94.51 = 70.28, P < 0.001 (see Figure 2 for details), with the probability judgments as dependent variable and the six levels of uncertainty as independent variable. The reaction times did not differ significantly between the six levels of uncertainty, F(2.59, 67.44) = 1.57, P = 0.21.

Model comparison and parameter estimation

To further explore how participants weighted the different types of information in belief updating, we compared the three different models described above according to their DIC values (see supplementary methods for details on model estimation and model comparison). The SM, which assumes a differential weighting of *social* as compared to *private information*, performed best ($\Delta_{\mathrm{DIC}_{\mathrm{FM\ minus\ SM}}}$ = 10.4; $\Delta_{\mathrm{DIC}_{\mathrm{BM\ minus\ SM}}}$ = 1590.4). This result was further supported by an analysis at the individual level: The Bayes factors favored the SM as compared to the FM for 24 of all 27 participants.

Figure 3 illustrates the difference in weighting of *social* and *private information* (SM) in belief updating. The weights given to *social information* (M_{social} - Figure 3A, left) were credibly smaller than the weights given to *private information* (M_{private} - Figure 3A, right). This is further illustrated by the contrast M_{private} - M_{social} (Figure 3B). Thus, during belief updating, participants substantially overweighted *private* as compared to *social information*. We also calculated the *information weighting index* on the basis of the estimated parameters of the SM for each participant. The *information weighting index* (Figure 3C) was significantly negatively correlated with the *conformity index*, Pearson's product moment correlation $r(25) = -.83$, $P < 0.001$, suggesting that the more people overweighted *private* as compared to *social information,* the less often they started a cascade in the trials of interest.

fMRI results

To investigate the neural processing of *social* and *private information* increasing uncertainty, we analyzed neural activity associated with belief updating.

General effects of belief updating

To correctly estimate the probability of choosing the better stock, a participant had to update her (prior) belief with every piece of information received (*social information I & II* and *private information*). Therefore, for the initial analysis we used a single parametric regressor that tracked the belief updating process independent of the social or private nature of the information (at $t = 2$ and 3 combined). Besides others, we found significant activity in fronto-parietal brain regions and in the precuneus during belief updating; that is, the activity of these regions increased with an increase in uncertainty (see Table 1, Figure 4 for further details).

Specific effects of belief updating by social or private information

Because our behavioral results indicated a differential processing of *private* and *social information*, we analyzed the two main belief updating stages (*social information II* and *private information*) independently. The left middle temporal gyrus/inferior parietal lobule was active during belief updating when subjects processed *social information II* (see Table 1) whereas activity of the anterior insula, the DLPFC, and the parietal cortex, besides others (see Table 1 and Figure 5), correlated with belief updating by *private information*.

Modulation of belief updating by individual differences in overweighting private information

The probability of an informational cascade starting depends on the differential weighting of *private* and *social information*. Therefore, we used the *information weighting index* to analyze how the process of belief updating (at $t = 3$) is modulated by inter-individual differences in information weighting. The regression analysis showed a positive correlation of the belief updating activity in the inferior frontal gyrus with the *information weighting index:* A similar positive correlation was observed in the anterior insula using a more liberal threshold ($P < 0.005$). Overall, the more participants overweighted *private* as compared to *social information,* the more active the inferior frontal gyrus/anterior insula were during belief updating of *private information* (Figure 5A and Table 2). An opposite effect was found in the parietal-temporal cortex: The more participants overweighted *private* as compared to *social information*, the less active the parietal-temporal cortex was during belief updating of *private information* (Figure 5B and Table 2).

DISCUSSION

By combining neurobiological, economic, and computational approaches, we were able to show that people who tend to overweight *private* as compared to *social information* show a decreased activity in the parietal-temporal cortex and an increased activity in the inferior frontal gyrus/anterior insula while updating their beliefs by *private information*. To our knowledge, this is the first study to illuminate the neural underpinnings of biased belief updating by *private information* – the cognitive process that is decisive for the emergence and stability of informational cascades.

Making an optimal decision when observing other people's decisions and receiving personal (private) information as represented by the informational cascades paradigm requires the integration of available s*ocial* and *private information* as described by the Bayesian solution. Deviations from the Bayesian solution (e.g. overweighting of *private information*) can influence subsequent decisions and therefore the occurrence of informational cascades. It is especially important for the theory of informational cascades to understand how the neural process of belief updating (of *private information*) is modulated by such deviations. The computational analysis of the behavioral data showed that subjects weighted *private* and *social information* differently: The majority of subjects (24 of 27 participants) overweighted *private* as compared to *social information*. This finding is consistent with recent research on informational cascades: A comprehensive meta-analysis by Weizsäcker (2010) showed that decision makers often overweight *private information* even in situations in which it would be optimal to follow others. The results of our behavioral control study (see supplementary results) indicate that subjects specifically overweight *private information,* which cannot alternatively be explained by an order-effect of overweighting recent information. Importantly, previous studies have shown that overweighting of *private information* strongly influences the emergence and stability of informational cascades (Bernardo and Welch, 2001; Nöth and Weber, 2003; Goeree *et al*., 2007). We also found a strong negative correlation between the individual tendency to make conforming decisions (*conformity index*) and overweighting of *private information* (*information weighting index*). This clearly indicates that overweighting of *private information* lowers the tendency to follow others and thereby lowers the probability that an informational cascade starts or continues.

Our fMRI results showed that an increase in uncertainty during belief updating by either s*ocial* or *private information* activated the parietal-temporal cortex – a region of the brain previously associated with number processing (Dehaene *et al*., 1998, 2003). Additionally, we found that an increase in uncertainty during belief updating by *private information* activated the DMPFC, bilateral anterior insula, and DLPFC – brain regions closely linked to decision risk (for a review, see Mohr et al., 2010). Furthermore, we demonstrated that stronger individual overweighting of *private information* positively correlated with activity in the inferior frontal gyrus/anterior insula and negatively with activity in the parietal-temporal cortex.

It has been shown that the inferior frontal gyrus is often co-active with the anterior insula (Paulus *et al*., 2003; Wright *et al*., 2012) and may constitute the so called “fronto-insular junction” (Craig, 2009). In the decision making under risk literature, activity of the inferior frontal gyrus has been related to higher risk aversion (Christopoulos *et al*., 2009), an increase in positive skewness (the chance of a better than average outcome is small) (Symmonds *et al*., 2011), an increase in the variance of an outcome (uncertainty) for risk-seeking individuals (Tobler *et al*., 2007), ambiguous *versus* non-ambiguous gambles, especially for ambiguity averse individuals (Bach *et al*., 2011), and increasing uncertainty (Huettel *et al*., 2005). Interestingly, a more posterior region within the inferior frontal gyrus was recently found to be more active the more improbable an event becomes as the result of a Bayesian updating process (d’Acremont *et al.*, 2013). Tracking of Bayesian posterior probabilities, however, has to be differentiated from belief updating of uncertainty as these are two different processes based on two different, but related, concepts (probability of occurrence with $0 \leq p \leq 1$ as compared to uncertainty with $0.5 \leq p \leq 1$). How belief updating leads to adjusted representations of posterior probabilities (i.e. the outcome of the belief updating process) is not yet known.

Activity in the anterior insula has been linked to risk anticipation (Mohr *et al*., 2010; Preuschoff *et al*., 2006), prediction of risk (Preuschoff *et al*., 2008), risk-aversion mistakes (Kuhnen and Knutson, 2005), intolerance of uncertainty (Simmons *et al*., 2008), risk during the selection of the potential behavioral responses (Huettel, 2006), and to the integration of subjective risk preference (Symmonds *et al*., 2011). Activity of the insular cortex has also been associated with the degree of harm avoidance (Paulus *et al*., 2003) and choice strategies that try to minimize losses (Venkatraman *et al*., 2009). Thus, we can speculate that the stronger uncertainty-related activity of the inferior frontal gyrus/anterior insula during the processing of *private information* conflicting with *social information* can overcome the effects of social conformity in subjective estimates of uncertainty.

However, according to the computational model (SM) overweighting of *private information* changes the posterior probability and thereby uncertainty. Thus, increased uncertainty could potentially explain increased activation of the inferior frontal gyrus/anterior insula in participants who strongly overweighted *private information*. To examine this explanation we determined whether overweighting of *private information* indeed increased uncertainty. The (un-)certainty measured as the average absolute difference between the posterior probability and a pure chance prediction of 0.5 across all trials was nearly the same for the SM with 0.2627 and the standard model (BM) with 0.2602. Therefore, overweighting of *private information* did not on average increase uncertainty and can be ruled out as an explanation for the increased activity of the inferior frontal gyrus/anterior insula. Instead, it appears plausible that people who are very sensitive to cues associated with uncertainty as reflected in increased activity of the inferior frontal gyrus/anterior insula tend to overweight *private information*. Overall, our results further support the important role of the anterior insula in the neural mechanism of social influence on human behavior (Berns *et al*., 2010; Campbell-Meiklejohn *et al*., 2010; Izuma and Adolphs, 2013; Klucharev *et al*., 2009).

The parietal-temporal cortex was active at all stages of belief updating (by *social* and *private information*). Importantly, activity of the parietal-temporal cortex was modulated by inter-individual differences in the weighting of *private information*: Stronger overweighting of *private information* was associated with decreased activity in the parietal-temporal cortex during the final stage of belief updating. Previous human and non-human studies consistently associated the parietal cortices with number processing (Dehaene *et al*., 1998, 2003) and with the resolution of uncertainty in tasks with limited knowledge about the correct action to take (Huettel *et al*., 2005, 2006; Kiani and Shadlen, 2009; Symmonds *et al*., 2011; Volz *et al*., 2003, 2004). Our results suggest that people with stronger numerical processing of *private information* in the parietal cortices are less biased towards *private information* and estimate uncertainty closer to the Bayesian optimal solution; however, this makes them more prone to start an informational cascade. Overall, we suggest a two-fold neural mechanism of overweighting of *private information* in informational cascades: (1) increased activity of the inferior frontal gyrus/anterior insula and (2) decreased activity in the parietal-temporal cortex. At a later stage during decision making, these two neural signals could be integrated via the direct anatomical connection between insula and posterior parietal cortex (Cavada and Goldman-Rakic, 1989). Further experiments are needed to explore this hypothesis.

We found a large overlap of activations evoked by increased uncertainty during belief updating by *private information* and during decision making (see supplementary fMRI results). In both time windows, we observed uncertainty-related activity of the DMPFC, anterior insula, parietal cortex and DLPFC. A meta-analysis by Mohr *et al*. (2010) showed that these brain regions are more strongly activated for decision risk as compared to anticipation risk. In our task, all relevant information was already available after the presentation of *private information*. Therefore, participants had the opportunity to form a decision (i.e., select a stock) before the response cue. Thus, in our task it is difficult to differentiate the neural effects related to belief updating and decision making at the last stages of a trial. Interestingly, in contrast to Stern *et al*. (2010), we did not find activity in the anterior cingulate cortex during belief updating (even when using a very low uncorrected threshold of 0.05). This discrepancy could be caused by the differences in the statistical analysis and/or design of the two studies. In contrast to our study, participants in the evidence accumulation task used by Stern *et al*. (2010) (1) rated uncertainty after each of the information cues, (2) received only *private information*, (3) received a feedback after every trial, and (4) had the opportunity to decline a decision. Thus, further studies are needed to clarify the exact role of the anterior cingulate cortex in belief updating. Additional studies will also help to generalize the observed mechanisms to different social environments.

Taken together, we show that *private information* conflicting with *social information* activates brain regions associated with risk and uncertainty. Furthermore, activity of the inferior frontal gyrus/anterior insula and the parietal-temporal cortex were modulated by inter-individual differences in the overweighting of *private information*. The behavioral results indicate that such inter-individual differences can influence the probability that a cascade starts. By and large, our results suggest a profound role of the uncertainty-related neural activity in the formation of informational cascades.

REFERENCES

Anderson, L.R., Holt, C.A. (1997). Information cascades in the laboratory. *American Economic Review,* 87, 847–62. Available at: http://www.jstor.org/stable/2951328

Bach, D.R., Hulme, O., Penny, W.D., Dolan, R.J. (2011). The known unknowns: neural representation of second-order uncertainty, and ambiguity. *Journal of Neuroscience, 31,* 4811–20. Available at: http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=3166851&tool=pmcentrez&rendertype=abstract

Banerjee, A.V. (1992). A simple model of herd behavior. *Quarterly Journal of Economics, 107*, 797–817. Available at: http://www.jstor.org/stable/2118364

Bernardo, A.E., Welch, I. (2001). On the evolution of overconfidence and entrepreneurs. *Journal of Economics & Management Strategy, 10,* 301–30. Available at: 10.1162/105864001316907964

Berns, G.S., Capra, C.M., Moore, S., Noussair, C. (2010). Neural mechanisms of the influence of popularity on adolescent ratings of music. *NeuroImage, 49,* 2687–96. Available at: http://www.sciencedirect.com/science/article/pii/S1053811909011410

Bikhchandani, S., Hirshleifer, D., Welch, I. (1992). A theory of fads, fashion, custom, and cultural change as informational cascades. *Journal of Political Economy, 100,* 992–1026. Available at: http://www.jstor.org/stable/2138632

Brett, M., Anton, J.L., Valabregue, R., Poline, J.B. (2002). Region of interest analysis using an SPM toolbox [abstract] presented at the 8th International Conference on Functional Mapping of the Human Brain. *Neuroimage, 16(2),* abstract 497

Campbell-Meiklejohn, D.K., Bach, D.R., Roepstorff, A., Dolan, R.J., Frith, C.D. (2010). How the opinion of others affects our valuation of objects. *Current Biology, 20,* 1165–70. Available at: http://linkinghub.elsevier.com/retrieve/pii/S0960982210005956

Cavada, C., Goldman-Rakic, P.S. (1989). Posterior parietal cortex in rhesus monkey: II. Evidence for segregated corticocortical networks linking sensory and limbic areas with the frontal lobe. *Journal of Comparative Neurology, 287*, 422–45. Available at: http://dx.doi.org/10.1002/cne.902870403

Chari, V.V., Kehoe, P.J. (2004). Financial crises as herds: overturning the critiques. *Journal of Economic Theory, 119,* 128–50. Available at: http://www.sciencedirect.com/science/article/pii/S0022053103002254

Christopoulos, G.I., Tobler, P.N., Bossaerts, P., Dolan, R.J., Schultz, W. (2009). Neural correlates of value, risk, and risk aversion contributing to decision making under risk. *Journal of Neuroscience, 29*, 12574–83. Available at: http://www.jneurosci.org/cgi/content/abstract/29/40/12574

Cialdini, R.B., Goldstein, N.J. (2004). Social influence: compliance and conformity. *Annual Review of Psychology, 55,* 591–621. Available at: http://www.annualreviews.org/doi/abs/10.1146/annurev.psych.55.090902.142015

Craig, A.D. (2009). How do you feel — now? The anterior insula and human awareness. Nature Reviews. *Neuroscience, 10,* 59–70. Available at: http://www.nature.com/nrn/journal/v10/n1/full/nrn2555.html

d'Acremont, M., Schultz, W., Bossaerts, P. (2013). The human brain encodes event frequencies while forming subjective beliefs. *Journal of Neuroscience, 33,* 10887–197. Available at: http://www.jneurosci.org/content/33/26/10887.abstract

Dehaene, S., Dehaene-Lambertz, G., Cohen, L. (1998). Abstract representations of numbers in the animal and human brain. *Trends in Neurosciences, 21,* 355–61. Available at: http://www.sciencedirect.com/science/article/pii/S0166223698012636

Dehaene, S., Piazza, M., Pinel, P., Cohen, L. (2003). Three parietal circuits for number processing. *Cognitive Neuropsychology, 20,* 487–506. Available at: http://www.tandfonline.com/doi/abs/10.1080/02643290244000239

Dieckmann, A., Rieskamp, J. (2007). The influence of information redundancy on probabilistic inferences. *Memory & Cognition, 35,* 1801–13. Available at: http://dx.doi.org/10.3758/BF03193511

Ellis, C.J., Fender, J. (2011). Information cascades and revolutionary regime transitions. *The Economic Journal, 121,* 763–92. Available at: http://dx.doi.org/10.1111/j.1468-0297.2010.02401.x

Goeree, J.K., Palfrey, T.R., Rogers, B.W., McKelvery, R.D. (2007). Self-correcting information cascades. *Review of Economic Studies, 74,* 733–62. Available at: http://search.ebscohost.com/login.aspx?direct=true&db=buh&AN=25378406&site=ehost-live

Huettel, S.A. (2006). Behavioral, but not reward, risk modulates activation of prefrontal, parietal, and insular cortices. *Cognitive, Affective, & Behavioral Neuroscience, 6,* 141–51. Available at: http://link.springer.com/article/10.3758%2FCABN.6.2.141

Huettel, S.A., Song, A.W., McCarthy, G. (2005). Decisions under uncertainty: Probabilistic context influences activation of prefrontal and parietal cortices. Jo*urnal of Neuroscience, 25,* 3304–11. Available at: http://www.jneurosci.org/content/25/13/3304.abstract

Huettel, S.A., Stowe, C.J., Gordon, E.M., Warner, B.T., Platt, M.L. (2006). Neural signatures of economic preferences for risk and ambiguity. *Neuron, 49,* 765–75. Available at: http://www.sciencedirect.com/science/article/pii/S089662730600078X

Hung, A.A., Plott, C.R. (2001). Information cascades: Replication and an extension to majority rule and conformity-rewarding institutions. *American Economic Review, 91,* 1508–20. Available at: http://www.jstor.org/stable/2677936

Izuma, K., Adolphs, R. (2013). Social manipulation of preference in the human brain. *Neuron, 78,* 563–73. Available at: http://www.sciencedirect.com/science/article/pii/S0896627313002705

Kiani, R., Shadlen, M.N. (2009). Representation of confidence associated with a decision by neurons in the parietal cortex. *Science, 324,* 759–64. Available at: http://www.sciencemag.org/content/324/5928/759

Klucharev, V., Hytönen, K., Rijpkema, M., Smidts, A., Fernández, G. (2009). Reinforcement Learning Signal Predicts Social Conformity. *Neuron, 61,* 140–51. Available at: http://www.sciencedirect.com/science/article/pii/S0896627308010209

Kuhnen, C.M., Knutson, B. (2005). The neural basis of financial risk taking. *Neuron, 47,* 763–70. Available at: http://www.sciencedirect.com/science/article/pii/S0896627305006574

Mohr PNC, Biele G, Heekeren HR (2010). Neural Processing of Risk. *Journal of Neuroscience,* 30, 6613–6619. Available at: http://www.jneurosci.org/content/30/19/6613.abstract

Nöth, M., Weber, M. (2003). Information Aggregation with random ordering: cascades and overconfidence. *The Economic Journal, 113,* 166–89. Available at: http://onlinelibrary.wiley.com/doi/10.1111/1468-0297.00091/abstract

Oberholzer-Gee, F. (2008). Nonemployment stigma as rational herding: A field experiment. *Journal of Economic Behavior & Organization, 65,* 30–40. Available at: http://www.sciencedirect.com/science/article/pii/S0167268106001119

Paulus, M.P., Rogalsky, C., Simmons, A., Feinstein, J.S., Stein, M.B. (2003). Increased activation in the right insula during risk-taking decision making is related to harm avoidance and neuroticism. *NeuroImage, 19,* 1439–48. Available at: http://www.sciencedirect.com/science/article/pii/S1053811903002519

Preuschoff, K., Bossaerts, P., Quartz, S.R. (2006). Neural differentiation of expected reward and risk in human subcortical structures. *Neuron, 51,* 381–90. Available at: http://www.sciencedirect.com/science/article/pii/S0896627306005046

Preuschoff, K., Quartz, S.R., Bossaerts, P. (2008). Human Insula activation reflects risk prediction errors as well as risk. *Journal of Neuroscience, 28,* 2745–52. Available at: http://www.jneurosci.org/content/28/11/2745

Raafat, R.M., Chater, N., Frith, C. (2009). Herding in humans. *Trends in Cognitive Sciences, 13,* 420–8. Available at: http://www.sciencedirect.com/science/article/pii/S1364661309001703

Selten, R. (1998). Axiomatic characterization of the quadratic scoring rule. *Experimental Economics, 1,* 43–62. Available at: http://dx.doi.org/10.1007/BF01426214

Simmons, A., Matthews, S.C., Paulus, M.P., Stein, M.B. (2008). Intolerance of uncertainty correlates with insula activation during affective ambiguity. *Neuroscience Letters, 430,* 92–97. Available at: http://www.sciencedirect.com/science/article/pii/S0304394007011275

Stern, E.R., Gonzalez, R., Welsh, R.C., Taylor, S.F. (2010). Updating beliefs for a decision: neural correlates of uncertainty and underconfidence. *Journal of Neuroscience, 30,* 8032–41. Available at: http://www.jneurosci.org/content/30/23/8032.abstract

Symmonds, M., Wright, N.D., Bach, D.R., Dolan, R.J. (2011). Deconstructing risk: Separable encoding of variance and skewness in the brain. *NeuroImage. 58,* 1139–49. Available at: http://www.sciencedirect.com/science/article/pii/S105381191100752X

Tobler, P.N., O'Doherty, J.P., Dolan, R.J., Schultz, W. (2007). Reward value coding distinct from risk attitude-related uncertainty coding in human reward systems. *Journal of Neurophysiology, 97,* 1621–32. Available at: http://jn.physiology.org/content/97/2/1621

Venkatraman, V., Payne, J.W., Bettman, J.R., Luce, M.F., Huettel, S.A. (2009). Separate neural mechanisms underlie choices and strategic preferences in risky decision making. *Neuron, 62,* 593–602. Available at: http://www.sciencedirect.com/science/article/pii/S0896627309002888

Vickery, T.J., Jiang, Y.V. (2009). Inferior parietal lobule supports decision making under uncertainty in humans. *Cerebral Cortex, 19,* 916–25. Available at: http://cercor.oxfordjournals.org/cgi/content/abstract/19/4/916

Volz, K.G., Schubotz, R.I., von Cramon, D.Y. (2003). Predicting events of varying probability: uncertainty investigated by fMRI. *NeuroImage, 19,* 271–80. Available at: http://www.sciencedirect.com/science/article/pii/S1053811903001228

Volz, K.G., Schubotz, R.I., von Cramon, D.Y. (2004). Why am I unsure? Internal and external attributions of uncertainty dissociated by fMRI. *NeuroImage, 21,* 848–57. Available at: http://www.sciencedirect.com/science/article/pii/S1053811903006797

Weizsäcker, G. (2010). Do We follow others when we should? A simple test of rational expectations. *American Economic Review, 100,* 2340–60. Available at: http://www.aeaweb.org/articles.php?doi=10.1257/aer.100.5.2340

Wright, N.D., Symmonds, M., Hodgson, K., Fitzgerald, T.H.B., Crawford, B., Dolan, R.J. (2012). Approach–avoidance processes contribute to dissociable impacts of risk and loss on choice. *Journal of Neuroscience, 32,* 7009–20. Available at: http://www.jneurosci.org/content/32/20/7009.abstract

Figure 1. Informational cascades task trial structure. The decisions of trader 1 (*social information I*) and trader 2 (*social information II*) were followed by a buying recommendation of a rating agency for one or the other stock (*private information*). At the end of every trial participants decided which stock (W or S) provided the higher revenue and indicated the probability of the correct outcome (probability judgment). The different windows were separated with fixation crosses (see “experimental design” section for details).

Figure 2. The effect of the different levels of uncertainty signaled by *social* and *private information* on participants’ probability judgments. An increase in objective certainty (x-axis) led to increased probability judgments (y-axis).

Note: the dotted line indicates the prediction of the normative Bayesian model (cf. Equation). The boxes range from the lower quartile to the upper quartile of the distribution. The black band in the middle of the box represents the median. The whiskers represent the minimum and the maximum of the distribution as long as these estimates are not further away from the median than ±1.5×IQR. Circles represent outliers.

Figure 3. Different weighting of *social* and *private* information (SM). (A) Marginal posterior distributions for the weight of the *social information* (M_{social}) and for the weight of the *private information* (M_{private}). (B) The contrast *private information* minus *social information* (M_{private} - M_{social}) indicates a strong difference of weighting of *social* and *private* information. (C) The distribution of the *information weighting index* shows that the majority of subjects overweight *private* as compared to *social information*.

Note: The 95% Highest Density Interval (95% HDI) spans 95% of the distribution. The vertical red line indicates hypothetical unbiased information weighting (i.e., equal weighting of *social* and *private information*).

Figure 4. Neural correlates of belief updating by *social* and *private information***.** Neural activity of the frontal and parietal cortices increased with increasing uncertainty of the decision.

Note: $P < 0.001$; cluster size = 20, uncorrected.

Figure 5. Inter-individual differences in belief updating by *private information*. Blue color indicates brain regions whose activity increased with increasing uncertainty during belief updating by *private information*. Results of the regression analysis (red boxes) represent activity of the subregions within the inferior frontal gyrus (A) and the parietal-temporal cortex (B) that was significantly correlated with overweighting of *private information* (*information weighting index*): The green color indicates a positive correlation, whereas the red color indicates a negative correlation. The two scatterplots display the average contrast estimates per subject within the respective cluster plotted against the information weighting index. The dashed red line displays a linear regression model.

Note: $P < 0.001$; cluster size = 0, uncorrected; brain regions in blue color – $P < 0.005$; cluster size = 0, uncorrected (explicit mask). Clusters are overlayed on a chi2better.nii.gz template provided by MRIcron (http://www.mccauslandcenter.sc.edu/mricro/mricron/).

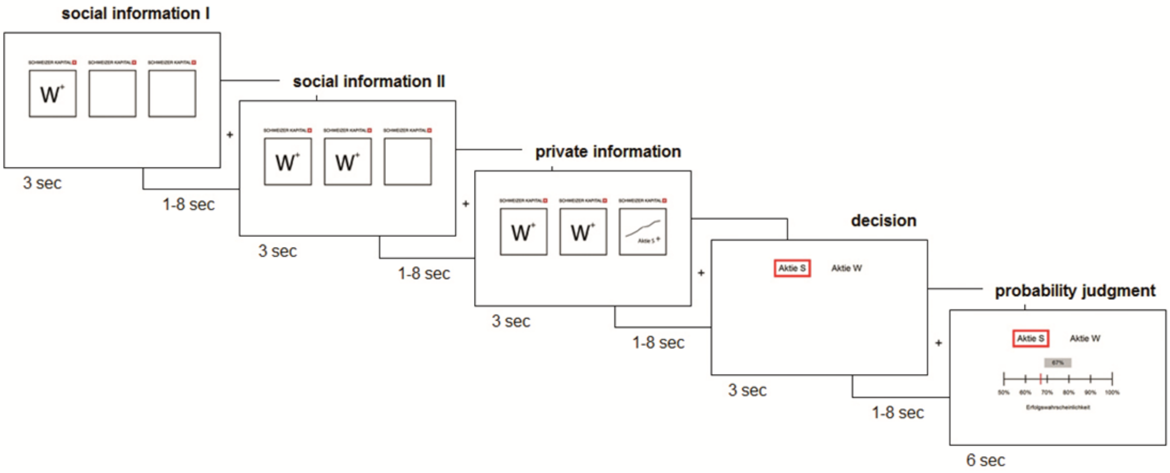
social information I
W+
3 sec
1-8 sec
social information II
W+
W+
3 sec
1-8 sec
private information
W+
W+
3 sec
1-8 sec
decision
Aktie S
Aktie W
3 sec
1-8 sec
probability judgment
Aktie S
Aktie W
6 sec

Figure 1.

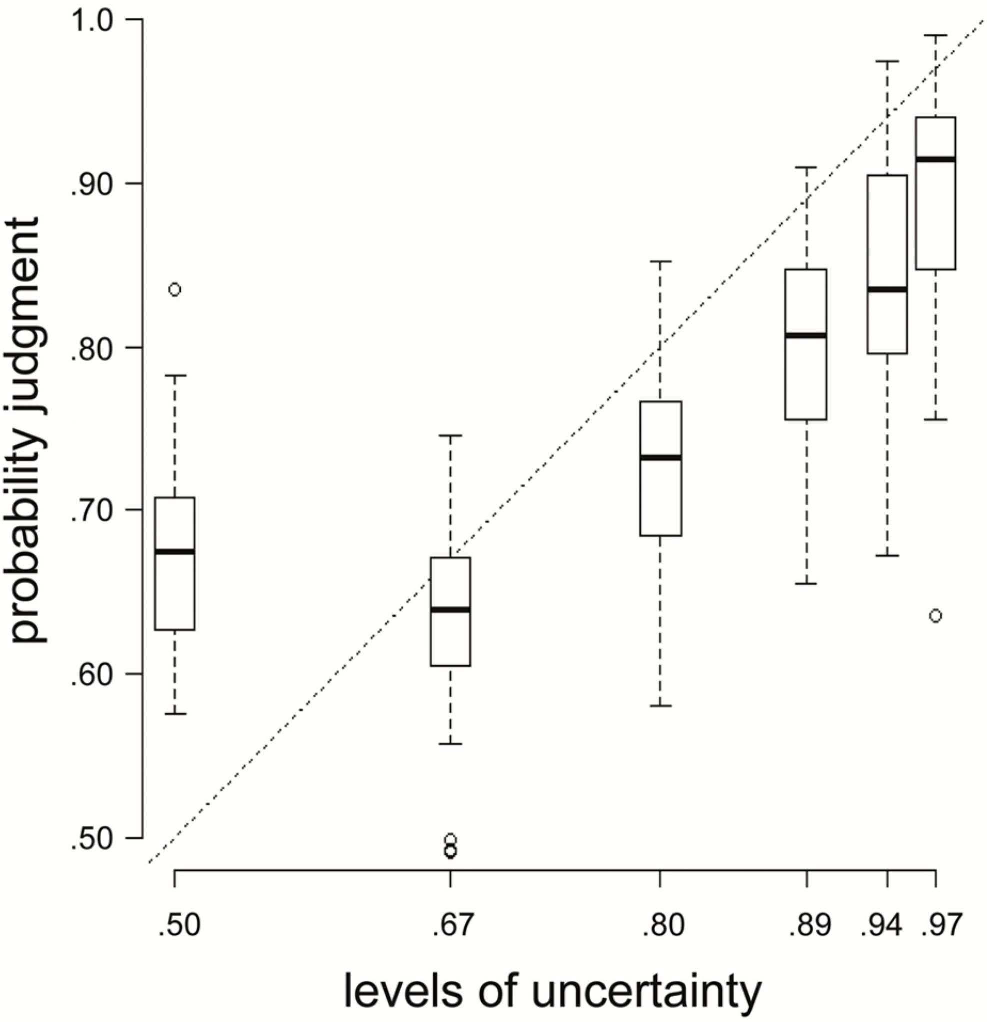
1.0
.90
.80
.70
.60
.50
probability judgment
.50
.67
.80
.89
.94
.97
levels of uncertainty

Figure 2.

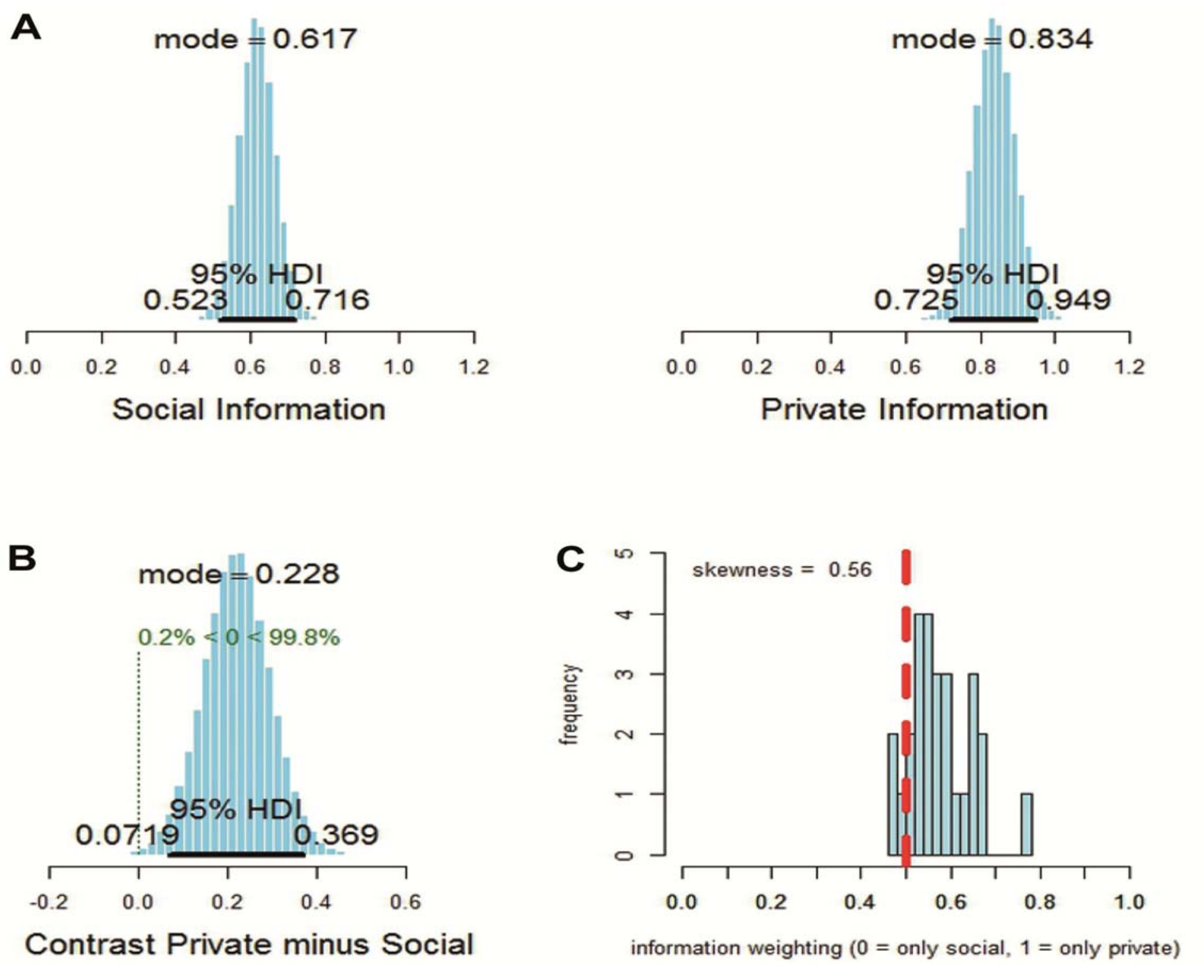
A
mode = 0.617
95% HDI
0.523
0.716
0.0 0.2 0.4 0.6 0.8 1.0 1.2
Social Information
mode = 0.834
95% HDI
0.725
0.949
0.0 0.2 0.4 0.6 0.8 1.0 1.2
Private Information
B
mode = 0.228
0.2% < 0 < 99.8%
95% HDI
0.0719
0.369
-0.2 0.0 0.2 0.4 0.6
Contrast Private minus Social
C
skewness = 0.56
frequency
0 1 2 3 4 5
0.0 0.2 0.4 0.6 0.8 1.0
information weighting (0 = only social, 1 = only private)

Figure 3.

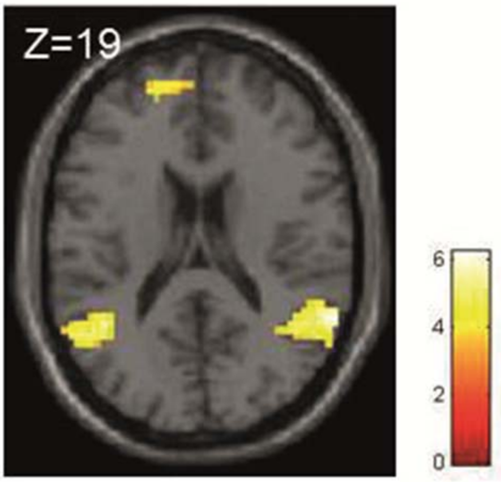

Figure 4.

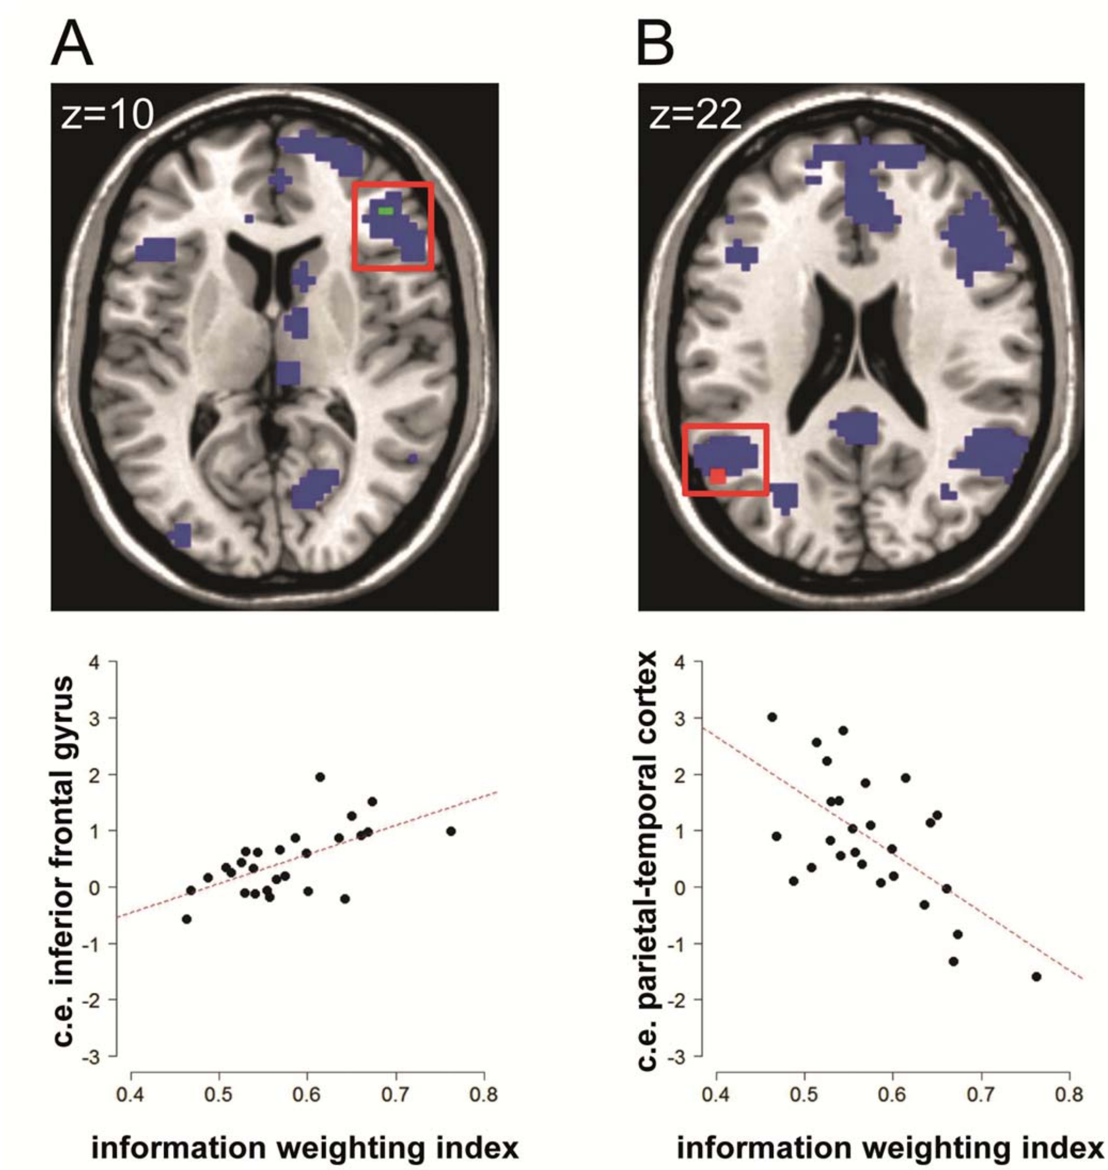
A
z=10
B
z=22
c.e. inferior frontal gyrus
c.e. parietal-temporal cortex
information weighting index
information weighting index

Figure 5.

Table 1. Neural correlates of belief updating

Contrast	Region	MNI centroid x	y	z	No. of voxel	Z value
General effects of belief updating (independent of *social* and *private information*)	Superior Temporal Gyrus/Inferior Parietal Cortex	63	-49	19	221	4.83
	Precuneus/Posterior Cingulate	3	-61	34	222	4.79
	Superior/Middle Frontal Gyrus (DLPFC)	-15	29	52	104	4.77
	Superior Temporal Gyrus/Inferior Parietal Cortex	-42	-61	28	229	4.54
	Superior/Middle Frontal Gyrus	21	26	46	59	3.99
	Superior/Medial Frontal Gyrus	-18	53	19	35	3.87
Belief updating by *social information*	Middle Temporal Gyrus	-42	-58	22	28	3.79
Belief updating by *private information*	Superior/Middle Frontal Gyrus (DLPFC)/DMPFC	48	32	19	1807	5.84
	Precuneus/Posterior Cingulate	6	-58	40	309	5.69
	Inferior Frontal Gyrus/Anterior Insula	48	41	-14	205	5.36
	Inferior Parietal Lobe	33	-64	40	524	4.83
	Inferior Parietal Lobe	-48	-64	43	372	4.52
	Middle Occipital Gyrus	27	-88	-5	145	4.45
	Middle Temporal Gyrus	42	-52	-11	161	4.41
	Cerebellum	-33	-73	-38	298	4.24
	Inferior Frontal Gyrus/Anterior Insula	-33	20	-2	108	4.17
	Middle/Inferior Frontal Gyrus	-39	41	-8	49	4.13
	Middle Occipital Gyrus	-36	-64	-11	120	4.07
	Parahippocampal Gyrus	21	-28	-11	20	3.95
	Dorsal Striatum	12	14	7	20	3.85

Table 2. Neural correlates of inter-individual differences in overweighting private information

Contrast	Region	MNI centroid x	y	z	No. of voxel	Z value
Positive correlation with *Information Weighting Index*	Inferior Frontal Gyrus	45	38	10	3 (11)	3.35
	Inferior Frontal Gyrus/Anterior Insula	39	17	-5	(3)	(2.90)
Negative correlation with *Information Weighting Index*	Middle Temporal Gyrus	-51	-64	22	8 (36)	3.46
	Midbrain	-3	-10	-11	1 (3)	3.45
	Middle Temporal Gyrus	-54	2	-23	1 (8)	3.38
	Middle Temporal Gyrus	-48	11	-29	1	3.35
	Midbrain	-6	-13	-8	1	3.24
	Middle Temporal Gyrus	-63	-31	-8	1 (20)	3.13
	Middle Temporal Gyrus	-51	2	-29	1	3.11
	Precuneus	-3	-52	40	1 (14)	3.11
	Middle Temporal Gyrus	-57	-31	-11	1	3.10
	Cerebellum	-33	-85	-38	(8)	(3.01)
	Middle Frontal Gyrus	-39	17	52	(3)	(2.86)
	Cerebellum	-15	-88	-38	(3)	(2.81)
	Medial Frontal Gyrus	-6	50	46	(1)	(2.64)
	Cerebellum	-18	-82	-29	(1)	(2.63)

Note: Z values in brackets are significant at $P < 0.005$ (uncorrected), whereas Z values without brackets represent results significant at $P < 0.001$ (uncorrected). The same logic is applied for the no. of voxels.

SUPPLEMENTARY MATERIALS

Neural Underpinnings of Informational Cascades: Brain Mechanisms of Social Influence On Belief Updating

Rafael E. Huber, Vasily Klucharev, and Jörg Rieskamp

I. Supplementary methods - Model estimation

To estimate the free parameters of the three computational models, we applied a Bayesian hierarchical approach (Kruschke, 2011) implemented with the OpenBugs software (Lunn *et al*., 2009) and the BRugs package (Thomas *et al*., 2006) in R (R Development Core Team, 2011). All three models provide a point estimate for the posterior probability that one stock is better than the other stock. To compare the model predictions $\hat{Y}$ (see Equation 3) with the observed probability judgments of the participants, we first transformed the observed probabilities using a logit transformation. We then assumed a normal distributed error (σ_n) around $\hat{Y}$ for each participant *n* as an additional free parameter. The model parameters for the n^{th} participant (β_{0n}, β_{tn} and σ_n) were sampled from group distributions, whereas the parameters of these group distributions were sampled from higher order distributions. In our hierarchical model, explicit prior assumptions were specified at the top of the hierarchical model only, as all the downstream parameters were connected to the overarching values. The model parameters of interest for the n^{th} individual (that is, β_{tn} and β_{0n}) were sampled from normal (group) distributions with means M_t and M_0 and precisions T_t and T_0 (where SD = $1/\sqrt{T}$). The means M_t and M_0 were sampled from normal hyperparameter distributions with a prior mean of μ_t = 1 and a precision τ_t = 0.01 for all M_t and μ_0 = 0 and τ_0 = 0.01 for M_0 (notice that the chosen prior means μ_t and μ_0 represent the normative solution). The precisions T_t and T_0 were sampled from gamma distributions with Shape = 0.1 and Rate =

0.1. As prior distribution for the error component σ_n we defined a gamma distribution with parameters S and R, which were also sampled from hyperparameter distributions (see Kruschke (2011, p.443) for a detailed description). For an efficient estimation process, we used a thinning factor of 100 and an initial burn-in of 10,000. All final Markov chains had a length of 100,000.

Model comparison

To compare the models we estimated the Deviance Information Criterion (DIC) for all three computational models. The DIC is especially suited to hierarchical models, as it takes the goodness-of-fit and the effective number of free parameters into account. The model with the lowest DIC should predict a replicate data set best (Spiegelhalter *et al*., 2002). Additionally, we compared the FM with the SM on the individual level via approximate Bayes factors based on the best fitting parameters (modes of the marginal posterior distributions) using the Bayesian Information Criterion (see Raftery, 1995; Wagenmakers, 2007).

Functional imaging data acquisition

Functional MRI was performed with ascending slice acquisition using a T2*-weighted echo-planar imaging sequence using a 3T Siemens Magnetom Verio whole-body MR unit equipped with a 12-channel head coil; 40 axial slices; volume repetition time (TR), 2.28 s; echo time (TE), 30 ms; 80° flip angle; slice thickness, 3.0 mm; field of view (FoV) read, 228 mm; slice matrix 76×76. For structural MRI, we acquired a T1-weighted MP-RAGE sequence (176 sagittal slices; volume TR, 2.0 s; TE, 3.37 ms; 8° flip angle; slice matrix 256×256; slice thickness, 1.0 mm; no gap; FoV, 256 mm). We preprocessed the fMRI data using SPM8 (Wellcome Trust Center for Neuroimaging, University College London). We applied a slice time correction using the middle image as reference. Preprocessing was continued with spatial realignment to correct for head movement. T1 images were then co-registered to the mean functional image created in the previous step. This image was segmented into grey matter, white matter, and cerebrospinal fluid (CSF). In a next step, the data were normalized according to the Montreal Neurological Institute (MNI) template and smoothed with a Gaussian smoothing kernel (FWHM = 8 mm). The start of the experimental paradigm was triggered by the 7^{th} scanner pulse to account for magnetization equilibration and previous scans were excluded from the final analysis.

II. Supplementary behavioural results - Control Study

The standard informational cascades paradigm implies a fixed order of *social* followed by *private information*. However, due to the fixed order of the presented information, it could have been that a different weight assigned to the *private informatio*n simply represented an order effect in which the last piece of information is given larger weight (e.g., Hogarth & Einhorn, 1992). Therefore, in order to examine whether indeed the last piece of information was given larger weight, we conducted a control study in which we had an additional condition in which only *private information* was presented. Seventeen participants (mean age = 21.6 years, ± 1.7 SD, 20-25 years, 6 females) participated in this additional behavioral study that consisted of 60 trials with 8 filler trials, 26 standard trials (*social information I*, *social information II* and *private information*; similar to the fMRI study) and 26 control trials (*private information I, private information II* and *private information III*). To explore a potential order effect, we compared the standard and control trials with each other by estimating the SM model. This enabled us to examine whether *private information* is weighted differently as compared to *social information* or whether simply the last piece of information is given larger weight than the preceding information.

The analysis of the standard trials replicated the behavioral results of the fMRI study: A clear trend towards overweighting of *private information* was observed, $M_{private}$ - M_{social}, mode = 0.165, 95% HDI = -0.0316 – 0.394. Importantly, the analysis of the control trials (consisting of *private information* only) did not show overweighting of the last *private information* at the end of the trial: The marginal posterior for $M_{private\ III}$ (*private information III*; mode = 0.748, 95% HDI = 0.615-0.873) could not be credibly differentiated from the marginal posterior for $M_{private\ I+II}$ (*private information I & II*, mode = 0.746, 95% HDI of 0.609-0.876), as indicated by the 95% HDI for the contrast $M_{private\ III}$ - $M_{private\ I+II}$, mode = -0.005, 95% HDI = -0.189-0.18. The results of the control study, in particular the control condition, showed that the last piece of information is not overweighted due to a recency effect and no order effect was observed. Thus, we can conclude that the larger weight given to *private information* as compared to *social information* in the fMRI study was due to the private *versus* social character of the information.

III. Supplementary fMRI results - The effect of subjective uncertainty during decision making

The probability judgments (i.e. subjective posterior probabilities) provided by the participants are a very direct measure of subjective uncertainty. Additionally, we analyzed the effect of subjective uncertainty on brain activity during decision making (decision time-window). At the end of each trial, participants made a probability judgment about their decision. We found that increased subjective uncertainty activated the bilateral fronto-parietal network, the left fronto-insular cortex and the dorsomedial prefrontal cortex (DMPFC) (Fig. S1 and Table S1). Thus, the brain areas involved into the belief updating by *private information* were also engaged into the final decision-making process.

REFERENCES

Hogarth, R.M., Einhorn, H.J. (1992). Order effects in belief updating: The belief-adjustment model. *Cognitive Psychology, 24,* 1–55. Available at: http://www.sciencedirect.com/science/article/pii/001002859290002J.

Kruschke, J.K. (2011). Doing Bayesian data analysis: A tutorial with R and BUGS. Burlington, MA: Academic Press.

Lunn, D., Spiegelhalter, D., Thomas, A., Best, N. (2009). The BUGS project: Evolution, critique and future directions. *Statistics in Medicine*, *28,* 3049–67. Available at: http://onlinelibrary.wiley.com/doi/10.1002/sim.3680/abstract.

Raftery, A.E. (1995). Bayesian model selection in social research. *Sociological Methodology, 25,* 111–63. Available at: http://search.ebscohost.com/login.aspx?direct=true&db=sih&AN=14956299&site=ehost-live

Spiegelhalter, D.J., Best, N.G., Carlin, B.P., Linde, A. van der (2002). Bayesian measures of model complexity and fit. *Journal of the Royal Statistical Society Series B (Statistical Methodology), 64,* 583–639. Available at: http://onlinelibrary.wiley.com/doi/10.1111/1467-9868.00353/abstract

Thomas, A., O'Hara, B., Ligges, U., Sturtz, S. (2006). Making BUGS open. *R News, 6*, 12–7. Available at: http://cran.r-project.org/doc/Rnews/.

Wagenmakers, E.-J. (2007). A practical solution to the pervasive problems of *p* values. *Psychonomic Bulletin & Review, 14,* 779–804. Available at: http://dx.doi.org/10.3758/BF03194105

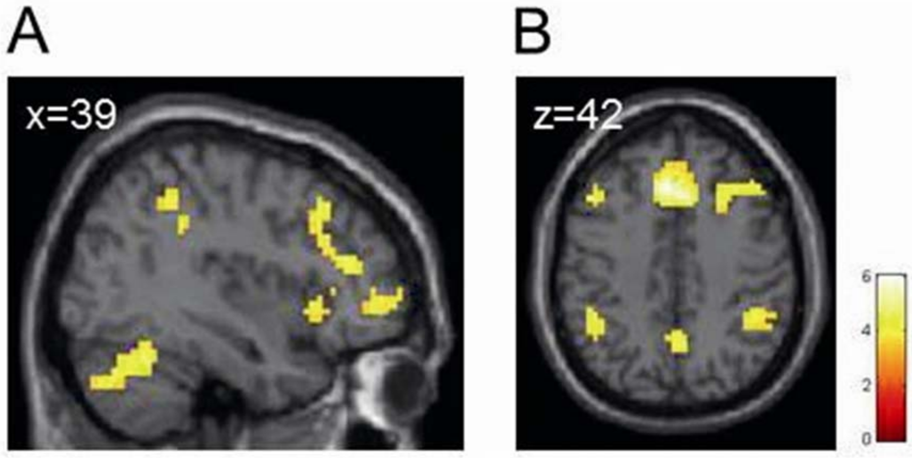

Figure S1. Neural correlates of subjective uncertainty during decision making. The decision-related activity of the anterior insular, parietal and frontal cortices increased with increasing subjective uncertainty.

Note: $p < 0.001$, cluster size = 20, uncorrected. Subjects indicated the subjective uncertainty of the decision at the end of each trial (probability judgment phase).

Table S1. Neural correlates of subjective uncertainty during decision making

Region	MNI centroid			No. of Voxel	Z value
	x	y	z		
Dorsomedial Prefrontal Cortex (DMPFC)	12	23	37	1368	4.74
Cerebellum	33	-61	-29	279	4.72
Inferior Frontal Gyrus / Precentral Gyrus	-48	41	1	644	4.65
Cerebellum	-27	-58	-32	647	4.51
Middle Temporal Gyrus	57	-40	1	110	4.31
Inferior Parietal Lobule	51	-46	22	146	4.24
Superior Temporal Gyrus	-54	-49	19	77	4.13
Superior Frontal Gyrus	-18	56	31	41	4.00
Thalamus	9	-13	13	24	3.89
Precuneus	3	-58	43	32	3.68
Inferior Parietal Lobule	-42	-52	43	76	3.63

Social Influences in Sequential Decision Making

Markus Schöbel, Jörg Rieskamp, and Rafael E. Huber

University of Basel, Department of Psychology, Basel, Switzerland

Author Note

This work was supported by the Swiss National Science Foundation (SNSF grant no. 100014-130352) to Jörg Rieskamp. Address all correspondence to: Markus Schöbel, Center for Economic Psychology, Department of Psychology, University of Basel, Missionsstrasse 62a, 4055 Basel, Switzerland, Phone: +41 (0)61 267 05 86, Fax: +41 (0)61 267 06 13, e-mail: m.schoebel@unibas.ch.

Abstract

The present work examines social influence on people's decisions in a sequential decision-making situation. The first experimental study using an information cascade paradigm shows that people infer information from others' decisions for making their own decisions. Following a cognitive modeling approach, our proposed social influence model shows that people overweight their own private information relative to the inferred social information. The second study examines the decision problem of Study 1 embedded in a medical decision-making problem. We test whether in the medical situation people do not only infer information from other's decision but also take other's authority into account. The social influence model illustrates an authority effect such that people overweight public information inferred from higher ranked persons as compared to equally ranked persons. Both studies shows how the social environment provides different sources of information that people integrate for making decisions.

Keywords: Social Influence, Conformity, Authority, Informational Social Influence, Information Cascade, Bayesian Analysis

Social Influences in Sequential Decision Making

Individuals often ignore their own opinion in favor of the opinions of others. Early experimental results of Asch (1951, 1956) and Sherif (1935) impressively illustrated how the judgments of others influence individuals' judgments. People sometimes follow the behavior of others even when they provide inaccurate information. The present article focuses on a decision-making problem in which several individuals sequentially make decisions and have the potential to influence each other. This situation has been studied by economists (e.g., Anderson & Holt, 1997; Bikhchandani, Hirshleifer, & Welch, 1992) who focused on conformity behavior due to the cognitive integration of socially inferred information improving individual decisions. In contrast, social psychologists have additionally emphasized conformity behavior, which is motivated by maintaining or building acceptance and belonging. Following a cognitive modeling approach, the goal of the present study is to examine to what extent individual decisions are affected by different types of social influence. Specifically, we are interested in how socially inferred information and normative expectations of an authority have an impact on individual decisions.

Imagine a physician confronted with the task of diagnosing a type of flu strain in a patient showing several symptoms. The symptoms speak in favor of Influenza A, but symptoms are only probabilistically related to flu strains. Thus the physician knows that her diagnosis will only be correct with a certain probability. Meanwhile she knows that her colleague has diagnosed a case of the relatively harmless Influenza C in the same patient. What should she do: Rely on the symptoms that she has observed or follow her colleague's judgment? If she follows her colleague's judgment this would be a typical case of conformity behavior, because she is disregarding the evidence the patient's symptoms provide. Can such a conformity decision be reasonable?

To explain why people conform it is helpful to distinguish two types of social influence: *normative social influence* and *informational social influence* (Deutsch & Gerard, 1955). Normative social influence describes behavior that has been driven by the desire to achieve a valued, coherent self-identity and to convey a particular impression to others (Chaiken, Wood, & Eagly, 1996). The influence is based on people's motivation to gain approval and avoid rejection by conforming with others' expectations. The physician's decision to conform may be motivated by the desire to avoid looking ridiculous in front of others because she was incapable of diagnosing the harmless Influenza C. In contrast, informational social influence arises from useful and valid information that another's opinion or behavior provides to improve a decision or judgment (Allen & Levine, 1971; Festinger, 1954). If, for instance, the physician's colleague was very experienced and potentially had additional information for a diagnosis, this informational influence would lead the physician to the correct inference that her colleague's diagnosis is very likely correct, making her own conforming decision the best she can do.

Dual-motive views of social influence have already been proposed in several domains, such as conformity research (Deutsch & Gerard, 1955; Insko, Drenan, Solomon, Smith, & Wade, 1983), group polarization research (Kaplan & Miller, 1983, 1987), and persuasion research (Wood, 2000). Criticism of such views has mainly focused on the problem of how the two types of influence can be separately measured and, consequently, how they interact (Allen, 1965; Levine & Russo, 1987; Tajfel, 1969). In many conformity studies individuals' behavior is examined under two conditions: In the public condition, individuals act under the surveillance of others, whereas in the private condition, responses are given anonymously. If behavior in the public condition differs from behavior in the private condition, this is usually attributed to salient beliefs of the person being socially influenced by the fact that others will positively evaluate his or her conformity behavior. Nevertheless, normative social influence cannot be excluded in the private condition. Social expectations of others can also emerge when their presence is imagined, so they hold across public and private contexts (Wood, 2000). Moreover, priming studies suggest that individuals' tendency to conform can even arise automatically, outside conscious awareness or voluntary control (Epley & Gilovich, 1999; Pendry & Carrick, 2001).

In sum, many social psychologists agree that conformity can result from informational and normative social influence. How the two types influence behavior is often difficult to measure, and whether and how they might work together is an even more complicated question. In the present study we examine a sequential decision-making task that helps us identify the different types of social influences on individual behavior. More specifically we examine decision making using the “information cascade paradigm” (e.g., Anderson & Holt, 1997; Bikhchandani, Hirshleifer, & Welch, 1992).

Information Cascades and Conformity Behavior

Bikhchandani et al. (1992) argued that people’s judgments, in principle, are based upon *private* and *public* information. For instance, based on a person’s own examination of a judgment situation, the person has access to information others have not obtained, which is private information. In addition, the person can consider information that is commonly available to everyone; this is public information. In a situation in which several individuals make the same decision sequentially, the decisions made by others preceding an individual’s own decision provide public information to that individual. An informational cascade occurs when it is optimal for an individual, having observed others’ preceding decisions, to follow the behavior of the preceding person, ignoring his or her own private information. Bikhchandani et al. (1992) showed that such decisions are rational when following a Bayesian analysis of the problem, which we demonstrate below.

Anderson and Holt (1997) examined whether information cascades actually occur. In their experiment, one of two urns was randomly selected by the experimenter. The two urns contained the same number of balls, but the composition of the balls’ color differs for the urns. For instance, both urns could contain three balls, with two white and one black ball for the first urn (Urn A) and two black and one white ball for the second urn (Urn B). The participants knew the compositions of the two urns but did not know which urn was randomly selected by the experimenter. Participants decided sequentially which of the two urns had been selected. Before making a decision, each participant drew one ball from the selected urn and observed its color, which was not revealed to the other participants (i.e., private information) and the drawn ball was afterwards put back into the urn. Thereafter, each participant publicly announced his or her decision. Thus, participants had private information, which was the color of the drawn ball from the chosen urn, and public information, which was the decisions of the preceding participants (but not their private signals). To make a correct prediction, participants could use both types of information.

More precisely, according to a Bayesian analysis of the problem, the posterior probability of an Urn A being selected could be determined by applying Bayes' theorem:

$$p(\mathrm{A} \mid n_\mathrm{a}, n_b) = \frac{p(n_\mathrm{a}, n_b \mid \mathrm{A})\, p(\mathrm{A})}{p(n_\mathrm{a}, n_\mathrm{b} \mid \mathrm{A})\, p(\mathrm{A}) + p(n_\mathrm{a}, n_\mathrm{b} \mid \mathrm{B})\, p(\mathrm{B})} \tag{1}$$

where $p(n_a, n_b \mid \mathrm{A})$ is the likelihood of observing the number n_a and n_b of "a" and "b" signals given Urn A was selected, where "a" speaks for Urn A and "b" speaks for Urn B. Signals are either obtained from private draws or inferred from public decisions of others. It is easier to determine the log odds of the posterior probability that Urn A was selected relative to the posterior probability that Urn B was selected. When assuming equal a priori probabilities with which the two urns are selected, the log odds are defined as

$$\ln \frac{p(\mathrm{A} \mid n_a, n_b)}{p(\mathrm{B} \mid n_a, n_b)} = \sum_{i=1}^{n_\mathrm{a}} \ln \frac{p(\mathrm{a} \mid \mathrm{A})}{p(\mathrm{a} \mid \mathrm{B})} + \sum_{i=1}^{n_\mathrm{b}} \ln \frac{p(\mathrm{b} \mid \mathrm{A})}{p(\mathrm{b} \mid \mathrm{B})}, \tag{2}$$

(for details see Appendix). When the log odds ratio is positive, then the posterior probability of Urn A being selected is larger than the posterior probability of Urn B being selected, whereas a negative ratio makes Urn B more likely to be selected. Under the assumption of equal priors and equal likelihoods of observing a or b signals, it can be easily seen with Equation 2 that solely the difference in the number of "a" and "b" signals is decisive (regardless of the absolute numbers of signals). For more details on the Bayesian solution to this problem see also Phillips and Edwards (1966), Grether (1980), Anderson and Holt (1997), or Hung and Plot (2001).

The following example illustrates the Bayesian analysis of the sequential decision problem. Suppose there are three people, named John, Jim, and Jack, facing the decision problem. John draws, unobserved by the others, the first ball and publicly decides for Urn A. After John's decision, Jim draws a ball and also decides for Urn A. Now it is Jack's turn. He draws a "b-ball," which indicates the selection of Urn B, but since John and Jim decided for Urn A, Jack infers that John has drawn an "a-ball," since he decided for Urn A. In addition, Jack infers that Jim also drew an a-ball, because if he had drawn a b-ball he probably would have decided for Urn B, to avoid being misled by a potential mistake of John.[1] Thus, Jack infers that two a-balls ($n_\mathrm{a} = 2$) and one b-ball ($n_\mathrm{b} = 1$) have been drawn and can calculate the log odds for Urn A:

$$\ln \frac{p(\mathrm{A} \mid n_\mathrm{a}, n_\mathrm{b})}{p(\mathrm{B} \mid n_\mathrm{a}, n_\mathrm{b})} = \sum_{i=1}^{n_\mathrm{a}} \ln \frac{2/3}{1/3} + \sum_{i=1}^{n_\mathrm{b}} \ln \frac{1/3}{2/3} = 0.69\,,$$

which are positive, so that Urn A should be selected despite the private signal supporting Urn B. Any subsequent decision makers should also follow the decision of the first and second decision makers, so that an information cascade emerges. If a fourth and fifth person drew b-balls it would be rational for them to decide for Urn A. Thus, although after the fifth person three b-balls and only two a-balls have been drawn, making Urn B the most likely selected urn, all individuals would be acting rationally by selecting Urn A according to a Bayesian analysis of the private and public information available to them.

Anderson and Holt (1997) observed a high proportion of individuals' decisions in line with the illustrated Bayesian updating process, which could be replicated by a multitude of empirical studies (e.g., Anderson, 2001; Hung & Plott, 2001; Kübler & Weizsäcker, 2004). However, compared to the Bayesian solution, participants in cascade experiments seem to overweight their private information relative to the public information (e.g., Bernardo & Welch, 2001; Goeree, Palfrey, Rogers, & McKelvey, 2007; Nöth & Weber, 2003). A recent meta-analysis by Weizsäcker (2010) comes to the overall conclusion that people often overweight their private information in comparison to public social information. However, we think that this conclusion needs to be limited to the artificial cascade paradigm examined. We think that people are often strongly influenced by other people's behavior in many real-life situations and thus overweight *social* relative to private information. Research illustrating the strong impact of social influences on behavior and decision making is widespread; for an overview, see for instance Cialdini & Goldstein (2004). In the present work we illustrate with the sequential decision-making paradigm described above how the impact of social influence can increase due to the social context in which it is embedded. Moreover, we follow a cognitive modeling approach to identify the importance people give to private as compared to social information.

Social influence model

To identify the importance people give to different sources of information we suggest a social influence model. For this model we modify Equation 2 by separating one component containing private from one component containing public information:

$$ln\frac{p(A|n_a,n_b)}{p(B|n_a,n_b)} =$$

$$\beta_{bias} + \beta_{soc}\left[\sum_{i\in a_{public}} f_i(a) + \sum_{i\in b_{public}} f_i(b)\right] + (2-\beta_{soc})\left[\sum_{i\in a_{private}} f_i(a) + \sum_{i\in b_{private}} f_i(b)\right] \quad (3)$$

where $f_i(a) = \ln\left[\frac{p(a|A)}{p(a|B)}\right]$, $f_i(b) = \ln\left[\frac{p(b|A)}{p(b|B)}\right]$, and where the index $i \in a_{private}, a_{public}$ denotes that the sum is reached over all private (public) "a" signals. The social importance parameter β_{soc} $(0 < \beta_{soc} < 2)$ specifies how much weight a person gives to the social as compared to the private information. In case of $\beta_{soc} > 1$ the decision maker overweights social information and in case of $\beta_{soc} < 1$ the decision maker overweights private information. The prior weight β_{bias} represents any initial bias towards one of the two choice options. Where $\beta_{soc} = 1$ and

$\beta_{bias} = 0$ the social influence model is equivalent to the Bayesian solution expressed by Equation 2. Note that the log odds of Equation 2 or 3 can be easily re-transformed into posterior probabilities by

$$p(A|n_a, n_b) = \frac{1}{1+e^{-ln\frac{p(A|n_a,n_b)}{p(B|n_a,n_b)}}}. \tag{4}$$

The larger the posterior probability for one option, the larger should be the probability that a person chooses this option. Accordingly we define the choice probability with which a person chooses an option as a function of the option's posterior probability of being correct:

$$p_{person}(A) = \frac{1}{1+e^{\theta\times(p(B|n_a,n_b)-p(A|n_a,n_b))}} \tag{5}$$

where θ (10>θ>0) represents a free sensitivity parameter that specifies how sensitive a person's response is to the different posterior probabilities. A large sensitivity parameter implies that the option with the larger posterior probability will be chosen with a larger probability.

In sum, the social influence model allows us to quantify the importance given to information inferred from others' decisions (public social information) relative to private information. By specifying the Bayesian solution as a special case of the model, we can test whether people deviate from the normative solution of probability theory.

In the second study, below, we manipulated the hierarchical rank of a previous decision maker to increase its social influence and the model allows us to test whether this manipulation affects the social influence. This was achieved by embedding the rather artificial cascade paradigm into a clinical decision-making context. Thereby, we draw on the authority principle, which states that people are willing to follow the suggestions of someone that they see as a legitimate authority (Cialdini, Bator & Guadagno, 1999; Milgram, 1974). The principle works within hierarchical relationships, which are asymmetrical in nature and involve the management of dominance "in ways that maximize the interests of the more dominant individual and limit harm to the less dominant individual" (Bugental, 2000, p.202). We understand the authority principle as a specific type of normative social influence, since it is based on the deference to authority norm, which is a prevailing norm in most organizations (Cialdini & Goldstein, 2004). However, manipulating normative social influence by confronting participants with a decision of a higher ranked person is a relatively weak induction of normative influence when compared to much more “pressurizing” homogenous majority opinion.

We examined the impact of social influence in two experiments by testing to what extent individual decisions are affected by social influences according to the following two hypotheses:

1. The *informational influence hypothesis* follows from a Bayesian view of information usage. This hypothesis states that people try to be as accurate in their judgments as they can be, efficiently inferring information from others’ behavior, and integrating the socially inferred information with their own private information to derive a decision. This decision can be the opposite of a decision that is reached from private information alone. Decision makers who behave in a manner consistent with the informational influence hypothesis will make decisions in line with the Bayesian model specified above (i.e., Equation 2). The social influence model allows us to test whether the decision maker weights all available information equally to make a decision, regardless of whether it is private or public information.

2. The *authority influence hypothesis* predicts that people’s behavior will also be influenced by the hierarchical status of other decision makers. In line with the authority principle, people will make decisions that conform to higher ranked others’ decisions more often, even if other available public information and one’s own private information suggest doing otherwise. Behavior that is consistent with the authority influence hypothesis should be better described by the social influence model, which allows the decision maker to give greater weight to the information that is inferred from the behavior of the higher ranked other person.

The aim of the following studies was to test these two hypotheses.

Study 1

The purpose of Study 1 was primarily to test the informational influence hypothesis. The experimental task was constructed in such a way as to minimize normative social influence on people's decisions, so that conformity behavior will largely express the informational social influence of others. If people's decisions are consistent with the Bayesian model, as suggested by Bikhchandani et al. (1992), this will indicate that individuals' decisions reflect a rational information integration process of privately and socially inferred information. In Study 1 we fit the social influence model to participants' decisions to see how and whether people's behavior deviates from the Bayesian solution.

The experimental task was similar to that used by Anderson and Holt (1997). However, to increase the experimental control, participants were not confronted with real urns from which balls were drawn. Instead they had to make judgments for a series of hypothetical scenarios (see Huck & Oechssler (2000) for a similar experimental procedure). This allowed us to systematically vary the information given to each participant. In contrast, in the experiment by Anderson and Holt participants had to announce their decisions to a group, so that normative social influence cannot be ruled out completely. In Study 1 participants were additionally asked to estimate the probability that their predictions were correct, so that we could compare it to the posterior probabilities derived by the Bayesian model (see Equation 2).

Method

Participants. A total of 40 students from different departments at the University of Basel participated in the 30-minute experiment. Participants received course credit or a book voucher worth 10 Swiss Francs. In addition, participants were informed that one of their decisions would be selected randomly, and if that decision were correct they would be rewarded with 2 Swiss Francs. If their corresponding confidence rating lay within the range of +-5% of the Bayesian solution they would receive an additional 2 Swiss Francs.

Procedure. Participants received a questionnaire with a description of the urn decision scenario, with two urns, each containing three balls, where Urn A had one black and two white balls and Urn B had two black and one white ball. Participants were instructed that one urn was randomly chosen at the beginning of the task by the experimenter and a maximum of four persons had the task of sequentially inferring which of the two urns was randomly chosen. They were told that four persons each sequentially drew one ball from the selected urn, which they replaced in the urn after they privately observed the ball’s color. Thereafter each person announced which urn he or she considered most likely to have been chosen. Thus each person knew the predicted urn of her or his predecessors (but not the color of their drawn balls). It was also explained that each person in the urn scenario had observed his or her predecessors’ decisions. Participants were told they should play the role of the person who made the last decision, in a total of 24 different scenarios.

After the situation description, participants received the 24 scenarios in a randomized order, in which the color of the ball that the last person had drawn and the decisions of the preceding persons were provided. The 24 scenarios presented 12 different decision tasks, where all possible combinations of up to four decision makers were specified. Decision sequences where participants were confronted with an unreasonable preceding decision (according to the Bayesian solution) were not included in our scenarios. The 12 decision tasks were presented in two different ways; that is, the decision sequences were mirrored in terms of the color of the balls and the decisions of the preceding persons. Thus, each participant decided twice on the same decision task. Participants were asked to predict for each scenario which urn (A or B) was most likely to have been randomly chosen by the experimenter. In addition, they had to judge the probability with which they thought their decision was correct (on a scale of 50–100%). Tables 1 and 2 together summarize 12 decision tasks with the corresponding posterior probabilities.

Results

We first analyzed whether participants' decisions were in line with the Bayesian solution. The fifth column of Table 1 shows the proportion of choices in line with the Bayesian solution (see Equation 1). For all tasks where the posterior probability was in favor of one alternative (Scenarios 1–9), 86.9% of all choices were consistent with the Bayesian prediction. In particular, when the Bayesian prediction was in favor of a participant’s private signal, 90.2% of all choices were consistent with the prediction. To determine whether information cascades occurred, Scenarios 6 and 8 are crucial. Here the Bayesian solution predicts that the private signal should be disregarded in favor of the previous decisions. A high degree of cascade behavior consistently occurred: Of all 160 choices, 120 (75.5%) were consistent with the Bayesian prediction.

In situations with posterior probabilities of p = .50 (Scenarios 10–12), private and public information cancel each other out. These scenarios allow us to test whether public social information has a stronger influence than private information. As shown in Table 2, in 79.9% of all choices participants decided in line with their private signal, thus participants gave more weight to their own information than to the public information[1]. In sum, the results show that participants used the information provided by others’ decisions in a way that is consistent with a Bayesian analysis of the decision problem, supporting the informational influence hypothesis.

To examine in more detail how much weight participants gave to public information relative to private information, we estimated the importance (β_{soc}), the bias (β_{bias}), and the sensitivity (θ) parameters of the social influence model on the basis of the observed data. We estimated the model by following a Bayesian approach for each participant (cf. Kruschke, 2010a, 2010b, 2011). This approach provides a posterior probability distribution of each of the model's free parameters. For each parameter, we first specified a prior distribution expressing the initial belief in every possible parameter. For the β_{bias} parameter we assumed a prior truncated normal distribution with a mean of zero and a standard deviation of 10, truncated at +1 and -1. For the social importance parameter β_{soc} we assumed a prior uniform distribution ranging from 0 to 2 (specified by a beta distribution). Likewise we assumed a uniform prior distribution ranging from 0 to 10 for the sensitivity parameter θ (specified by a beta distribution). According to the Bayesian approach, the prior distributions are then updated on the basis of the data and the model's likelihood function (i.e., Equation 5). Technically we relied on *JAGS* (Plummer, 2003) through the *rjags* interface in R (R Development Core Team, 2011). For the sampler we chose a thinning factor of 100 (to minimize autocorrelation) and an initial burn-in of 10000 (to produce more representative samples from the posterior). The final Markov chains had a net length of approximately 50,000. Group estimates for the parameters of the model were derived by averaging the posterior distributions of all participants (by averaging the results of the Markov chains). The derived distributions of the means can be used to calculate summary statistics (e.g., median, 95% highest density interval; HDI[2], etc.).

For the social influence model the median estimated sensitivity parameter was θ = 6.08 (*95% HDI* = 5.57 – 6.58), which implies that participants reacted rather sensitively to the different posterior probabilities. For instance, with a value of 6.08 for the sensitivity parameter, Urn A will be chosen with a probability of .89 given a posterior probability of .67 for Urn A. For β_{bias} the estimated median parameter value was -0.12 (95% HDI = -0.22 – -0.01), which indicates a slight tendency to favor Urn B a priori. The median importance parameter given to the public information was β_{soc} = 0.78 (*95% HDI* = 0.71 – 0.86), which shows that participants weighted public information less strongly than private information. When contrasting the weight given to private information (2- β_{soc}) with the weight given to public information (β_{soc}) a median positive difference of 0.44 results (*95% HDI* = 0.28 –

0.59), illustrating overweighting of private information. In sum, the analysis shows that participants overweight private as compared to public information—inconsistent with the Bayesian model that weights all information equally.

In addition to making choices between the two urns the participants had to judge the probability that their choices were correct. The probability judgments, reported in the last columns of Tables 1 and 2, did not match the Bayesian posterior probabilities. Whereas the average probability judgment of .59 was higher than the posterior probability of .50 in Scenarios 10-12, for scenarios with a posterior probability of .67, .80, and .89 the average probability judgments of .61, .69, and .74, respectively, were lower. These results appear similar to the standard conservatism phenomena reported in the early literature on probability judgments (Edwards, 1968), according to which people tend to give less moderate probability judgments. However, our social influence model might give an alternative explanation for these deviations. The social influence model, which we propose, predicts the probability with which people will select one or the other option (see Equation 5). These predictions follow from the models' predicted *subjective* posterior probabilities that one or the other option is correct. Therefore people's confidence judgments can also be compared to these subjective posterior probabilities that the model predicts. Importantly, the model was estimated solely on the basis of participants' choices ignoring their confidence judgments. Therefore predicting participants' confidence judgments represents a strong generalization test of the social influence model.

Figure 1A shows that the model predicts the observed confidence judgments very accurately. Importantly, the model also predicts overweighting of small probabilities and underweighting of large probabilities. For instance, the model correctly predicts a confidence level of 61% compared to people's observed confidence levels of 59% in situations in which the normative Bayesian account predicts a posterior probability of 50%. Similarly, for situations with a normative posterior probability of 89%, the model predicts a confidence judgment of 83% compared to the empirically observed confidence judgment of 75%. Thus, the social influence model can predict the observed deviations of people's confidence judgments from the normative account. According to the social influence model these deviations result from overweighting individual as compared to public information. For instance, in the normative indifference situation with posterior probabilities of 50% overweighting private information leads to increased confidence, whereas in situations with normative high posterior probabilities overweighting private information leads to more moderate confidence levels.

Discussion of Study 1

Study 1 shows that people decided to conform and go against their own private signal depending on whether the posterior probabilities spoke for or against their private signal. In situations where private and public information cancelled each other out, the private information was preferred over public information. These results suggest that private information and socially inferred information are cognitively integrated. Furthermore, the results replicate Anderson and Holt’s (1997) findings, where the participants made real draws from the urns. Our hypothetical scenarios have the advantage of maximizing experimental control. For instance, the scenarios minimize potential normative social influences of other people present in a public setting. Therefore, our results illustrate the impact of informational social influence leading to conformity behavior. The results of the social influence model show that participants overweight private as compared to public information, contrary to equal weighing of the Bayesian model. Likewise, participants’ probability judgments do not correspond to the Bayesian solution. These deviations from the Bayesian posterior probabilities could be explained by the social influence model. According to the social influence model people overweight their private information as compared to social information, which on average in the tested situations lead to more moderate confidence judgments. Importantly the model predicts these deviations from the Bayesian account without being fitted to the observed confidence judgments.

Study 2

The purpose of Study 2 was to investigate decision making in a real-life situation in which both informational and authority influences may affect people's decisions. Therefore, in Study 2 the decision problem of Study 1 was embedded in a medical decision-making context. Participants had to take on the role of an assistant physician who had to diagnose, on the basis of particular symptoms, from which of two diseases a patient was suffering. The task was analogous to that of Study 1: The assistant physicians had information about others' decisions; here, the previously made diagnoses of other physicians recorded in the patient's record. The other physicians' decisions were often not supported by the private information available to the assistant physician. Again, these decisions represent informational social influence to the assistant physician. To examine social influences of the hierarchical status of preceding decision makers, the cascade paradigm offers the opportunity to control the strength of authority influences by varying the hierarchical ranking of the preceding decision makers. At the same time, we can control the strength of informational influences by determining the validity of available information that the decision makers in a sequence draw on. In the following, we explain how we manipulate both types of social influences to examine the influence of authority relative to informational social influences.

To manipulate authority influence, the hierarchical ranking of the influence source was varied: The preceding decisions were made either by a colleague (another assistant physician) with the same hierarchical ranking or by a supervisor (the medical director) with a higher hierarchical ranking. This manipulation varied the strength of the authority influence by focusing on the legitimate power of previous decision makers in relation to the assigned hierarchical ranking of the participant's role. Although our participants did not expect any negative consequences when deciding against the diagnosis of the medical director, we argue that the tendency to conform should emerge as a result of the perceived hierarchical status difference in line with priming studies on conformity (Epley & Gilovich, 1999; Pendry & Carrik, 2001). To control the strength of informational social influences, participants were told that the average accuracy of the assistant physicians' and the medical director's diagnoses on the specific decision problem was the same. This allowed us to test the informational influence hypothesis and the authority influence hypothesis within the same task.

In Study 2, 40 scenarios were employed in which participants were confronted with the same 12 decision tasks of Study 1. In order to test our hypotheses, we created all possible variations of the same decision task in terms of varying hierarchical rankings of the previous decision makers. More specifically, 40 scenarios for all possible decision sequences for up to four decision makers were created, in which the medical director and assistant physicians decide at all positions in the decision sequence with corresponding diagnoses. Again, we excluded scenarios with unreasonable preceding decisions (according to a Bayesian analysis), e.g., scenarios where two decisions favoring the same diagnosis are followed by an opposed decision. In sum, we created four scenarios with one previous decision maker (one with the assistant physician and one with the medical director as previous decision makers favoring or opposing participants' private information), 12 scenarios with two previous decision makers and 24 with three previous decision makers (see Tables 3–6 in Column 1 and 2 for the scenarios used in Study 2).

In a next step, we structured the scenarios according to the corresponding Bayesian predictions, resulting in four groups of scenarios (scenarios with a posterior probability of .50, .67, .80, and .89; see Tables 3–6). This study design allowed us to test both social influence hypotheses. According to the informational influence hypothesis, we should obtain no differences in participants' decision making and probability judgments in (the four groups of) scenarios where the Bayesian solution is the same. However, according to the authority influence hypothesis, participants' decisions should vary depending on (a) whether the decision of the higher ranked decision maker (the medical director) supports or speaks against participants' privately held information and (b) whether the medical director is one of the preceding decision makers or not. Due to the same informational value of previous decisions, independently of the hierarchical status of preceding decision makers, changes in participants' decision making and probability judgments within a scenario group (i.e., a group of scenarios with the same Bayesian solution) could be traced back to the impact of the hierarchical status of previous decision makers. Therefore, we calculated the average proportion of participants' decisions in favor of their private information for the following three types of scenarios (within each of the four groups of scenarios i.e., of scenarios with a posterior probability of .50, .67, .80, and .89) (see Table 3-6):

1. Scenarios in which only assistant physicians are the preceding decision makers (baseline condition)
2. Scenarios in which the medical director's decision supports participants' private information
3. Scenarios in which the medical director's decision speaks against participants' private information

According to the authority hypothesis, we predict that (a) participants should decide more strongly according to their private information (and should be more confident) when the medical director supports it relative to decisions in the baseline condition (i.e., comparing scenarios b to a); and (b) participants should decide less according to their private information (and should be less confident) when the medical director's decision speaks against it relative to decisions in scenarios of the baseline condition (i.e., comparing scenarios c to a).

Method

Participants. A total of 40 students from different departments at the University of Basel participated in the experiment, which took approximately one hour. Participants received a course credit or a book voucher worth 10 Swiss Francs. In addition, participants were informed that one of their diagnoses would be selected randomly, and if that diagnosis were correct according to the Bayesian solution they would be rewarded with 2 Swiss Francs. If their corresponding confidence rating lay within the range of +-5% of the Bayesian solution they would receive an additional 2 Swiss Francs.

Procedure. First, participants received a description of a hypothetical situation in a hospital. They were asked to imagine themselves in the position of an assistant physician who had to make a decision concerning a patient's disease. Participants were told about two possible diseases, which were a priori equally likely: sigma diverticulitis and appendicitis. Both diseases were probabilistically related to two independently occurring symptoms. Participants were informed that the patient suffered from one of the two symptoms; this constituted the private information of the participant. The first symptom, regurgitation, was more often observed when patients suffered from sigma diverticulitis; that is, the conditional probability of observing the symptom when the patient suffered from the disease was *p*(regurgitation|sigma diverticulitis) = .67, whereas the conditional probability of observing the symptom when the patient suffered from appendicitis was *p*(regurgitation|appendicitis) = .33. The second symptom, twinges in the left underbelly, was more often observed when patients suffered from appendicitis; that is, *p*(twinges in the left underbelly|appendicitis) = .67, whereas the symptom was less often observed when patients suffered from sigma diverticulitis; that is, *p*(twinges in the left underbelly|sigma diverticulitis) = .33.

In addition, the scenarios provided public information concerning the previous diagnoses made by other assistant physicians and/or the medical director, which were recorded in the patient's record. Participants were informed about the average accuracy of the assistant physician and the medical director when making an independent diagnosis, that is, a diagnosis without knowing other physicians' diagnoses. Participants were told that an independent diagnosis of the assistant physician and the medical director was correct in 2 out of 3 cases (p = .67). Thus, the decisions of all preceding decision makers (independently of their hierarchical rank) have the same validity of being correct.

After the initial situation was described, 40 decision scenarios were given to the participants in a randomized order. The 40 scenarios provided the participants with the symptom of the patient and the previous diagnoses. Tables 3–6 summarize the 40 decision scenarios with the corresponding posterior probabilities. For each scenario participants were asked to predict which disease (appendicitis or sigma diverticulitis) the patient had developed. In addition, they were asked to judge the probability with which they thought their diagnosis would be correct (on a scale of 50–100%).

Results

The purpose of Study 2 was to examine individuals' decision making in relation to the predictions of the informational and the authority influence hypotheses. We broke down our analysis into three parts: First, we present the results of testing the informational influence hypothesis. Next, we describe the results of examining the authority influence hypothesis. Finally, we fit the observed decisions with the social influence model describing the interplay between informational and authority influences.

Informational social influences. To examine whether participants behaved according to the Bayesian analysis of the decision problem, we first analyzed their decisions. The fifth column of Tables 3–5 shows the proportion of participants who made choices in line with the posterior probabilities derived from the Bayesian analysis (see Equation 2). For all scenarios in which the posterior probability was in favor of one disease (Scenarios 1–30), 92.0% of all choices were consistent with the Bayesian prediction. In particular, when the Bayesian prediction was in favor of a participant's private information, 95.1% of all choices were consistent with the prediction. To determine if informational cascades occurred, Scenarios 3, 4 and 9–13 are crucial (Table 3). Here the Bayesian solution predicts that the private signal should be ignored in favor of the previous decisions. Consistently, a high degree of cascade behavior occurred: Of all 280 decisions, 230 (82.1%) were consistent with the Bayesian prediction.

In situations with posterior probabilities of $p = .50$ (Scenarios 31–40), private and public information cancel each other out. Similarly to Study 1, these scenarios allow us to test whether public information has a stronger influence than private information. As shown in Table 6, in only 49.7% of all diagnoses did participants decide in line with their private signal. To explain this result it is important to examine the influence of the authority influence presented in the next section. Overall, the results show that participants used information provided by others' decisions consistent with a Bayesian analysis of the problem, supporting the informational influence hypothesis.

Authority influences. In order to examine the authority influence on participants' decisions, we first analyzed whether participants in general decided against or with the diagnosis of the medical director. We had 1119 diagnosis decisions in scenarios where the medical director was one of the preceding decision makers. Of these, 838 (74.89%) were in line with the diagnosis of the medical director. However, to evaluate the impact of authority, it is crucial to focus on participants' diagnoses and probability judgments with regard to (1) the Bayesian prediction of each decision scenario and (2) the comparison of scenarios with and without the medical director as preceding decision maker supporting or disapproving participants' private information. Therefore, we draw on four scenario groups (see Tables 3–6), in which each scenario had the same posterior probability of one disease (Scenarios 1–30) or the posterior probabilities predict an indifference situation (Scenarios 31–40).

The authority hypothesis predicts that (a) participants should decide more strongly according to their private information (and should be more confident) when the medical director's decision supports their private information relative to decisions in scenarios of the baseline condition where the medical director is not one of the preceding decision makers. Likewise, (b) participants should decide less according to their private information (and should be less confident) when the medical director's decision speaks against their private information relative to decisions in scenarios of the baseline condition.

We began with scenarios for which the posterior probability of one disease according to a Bayesian analysis is .67 (see Table 3). The average proportion of participants' decisions favoring the private information is higher in scenarios where the medical director supports participants' private information compared to the baseline scenarios where the medical director is not one of the previous decision makers (z = -5.12, p = .001 according to a *Wilcoxon signed-rank test*). Moreover, we found a lower average proportion of decisions according to private information in scenarios where the medical director decided against participants' private information compared to the decisions at the baseline scenarios (z = -4.85, p = .001). Participants' probability judgments with M = .69 were higher in scenarios where the medical director's decision supported participants' private information compared to the baseline condition with M = .65; (t(39) = -2.54, p = .015). However, the probability judgments for scenarios where the medical director's decision speaks against participants' private information with M = .66 were not different to the probability judgments for the baseline scenarios with M = .65, (p = .07).

Next, we present the results of comparing participants' decisions in scenarios for which the posterior probability of one disease according to a Bayesian analysis is .80 (see Table 4). We found no significant difference of the average proportions of decisions according to the private information between scenarios where the medical director's decision favors the private information and the baseline scenarios (p = .65). However, participants decided less often according to their private information in scenarios where the decision of the medical director spoke against their private information compared to their decisions in the baseline scenarios (z = -2.23, p = .026), supporting our authority influence hypothesis. No significant differences in the probability judgments could be observed between scenarios where the medical director's decision favored the private information and the baseline scenarios. However, participants' decisions against the medical director's decisions showed a significantly lower confidence (*M* = .67) compared to the confidence judgments in the baseline scenarios (*M* = .77) (t(39) = 4.36, *p* = .001).

In scenarios in which the posterior probability of one disease is .89 (Table 5), we found no significant differences in participants' average proportion of decisions in line with their private information between scenarios where the medical director's decision corresponds to participants' private signal and the baseline scenarios (*p* = .32) whereas their probability judgments significantly differed between both conditions (t(39) = -3.55, *p* = .001) in the direction of a higher confidence for decisions which correspond with the medical director's decision.

Lastly, we analyzed decisions and probability judgments in scenarios where the posterior probabilities for both diseases were the same, with .50 predicting indifference for the diagnoses (see Table 6). We found no significant differences between scenarios where the decision of the medical director favors participants' information and decisions made in the baseline scenarios. Consistent with the authority influence hypothesis, we found a significantly lower average proportion of participants' decisions in line with their private information in scenarios with participants' private information opposite to the medical director's decision compared to the baseline scenarios where the medical director is not one of the preceding decision makers (*z* = -3.42, *p* = .001). Participants' probability judgments were significantly higher in scenarios where the medical director supports the private information (*M* = .69) compared to the probability judgments at the baseline scenarios (*M* = .63) (t(39) =-5.18, *p* = .001). Moreover, the probability judgments in scenarios where the decisions of the medical director speak against participants' private information (*M* = .65) were significantly higher compared to the probability judgments at the baseline scenarios (t(39) = -2.54, *p* = .015).

In sum, we found strong empirical evidence for our authority hypothesis when comparing participants' decisions in scenarios without the medical director as preceding decision maker (baseline scenarios) with scenarios, in which the medical directors' decision contradicts participants' private information. Here, the average proportion of decisions according to private information indicates a consistent tendency to follow authority influences. The analysis of the impact of authority influences supporting participants' private information provided evidence that for scenarios with a posterior probability of .67 participants more often decided according to their private information (compared to their decisions at the baseline scenarios), whereas this influence was not observed for scenarios with posteriors of .80 and .89. This could be due to a ceiling effect, because for the scenarios with high posterior probabilities we had already observed high proportions of decisions in line with private information in the baseline scenarios. However, the probability judgments were consistently higher in scenarios with a supporting decision of the medical director compared to the probability judgments in the baseline scenarios illustrating an authority influence.

The social influence model. Finally, we estimated the social influence model on the basis of participants' decisions. The goal in Study 2 was to distinguish informational from authority influence. Therefore, we decomposed the public information component within Equation 3 into two components instead of only one; one referring to information from higher ranked decision makers and one referring to information from equally ranked decision makers, providing:

$$ln\frac{p(A|n_A,n_B)}{p(B|n_A,n_B)} = \beta_{bias} + \beta_{higher}\left[\sum_{i\in a_{public\,higher\,ranked}} f_i(a) + \sum_{i\in b_{public\,higher\,ranked}} f_i(b)\right] + \beta_{equal}\left[\sum_{i\in a_{public\,equal\,ranked}} f_i(a) + \sum_{i\in b_{public\,equal\,ranked}} f_i(b)\right] + \left(3 - \beta_{higher} - \beta_{equal}\right)\left[\sum_{i\in a_{private}} f_i(a) + \sum_{i\in b_{private}} f_i(b)\right] \quad (6)$$

where β_{higher} refers to the importance given to the information derived from the decisions of the higher ranked medical director and β_{equal} refers to the importance given to the information derived from the decisions of the equally ranked assistant physician. In the case of β_{higher} = 1 and β_{equal} = 1 the social influence model specified by Equation 6 is identical with the pure Bayesian model (see Equation 2). To estimate the four free parameters ($\beta_{bias}, \beta_{higher}, \beta_{equal}$ and θ - see Equation 6) of the social influence model for every participant in Study 2 we applied the same Bayesian approach as used in Study 1 (except that we used a $precision$ (SD= $1/\sqrt{precision}$) of 0.01 instead of 0.1 for the prior distribution of β_{bias})[3].

The median estimated sensitivity parameter for the social influence model in Study 2 was θ = 7.37 (*95% HDI* = 6.9 – 7.85), thus just a little higher than in Study 1. The median parameter estimate for β_{bias} was 0.06 (95% HDI = -0.02 – 0.13), indicating no a priori bias toward one of the two decision options. The median importance parameter for $\beta_{public\ higher\ rank}$ was 1.12 (*95% HDI* = 1.05 – 1.19), which was higher as compared to the median importance parameter $\beta_{public\ equal\ rank}$ = 0.85 (*95% HDI* = 0.78 – 0.91). The contrast between the two parameters $\beta_{public\ higher\ rank}$ - $\beta_{public\ equal\ rank}$ was positive with a median difference of 0.27 (*95% HDI* = 0.15 – 0.39). The median weight for the private information of 1.03 (*95% CI* = 0.97 – 1.10) shows that participants gave more weight to private information than public information derived from the decisions of the equally ranked doctors ($Median_{difference}$ = 0.19 (*95% HDI* = 0.08 – 0.30). However, private information was not treated differently when compared to the importance given to information derived from the decisions of the higher ranked doctors ($Median_{difference}$ = -0.09 (*95% HDI* = -0.21 – 0.04). Therefore, the results of the social influence model show that people give greater weight to public information derived from higher ranked individuals than public information derived from equally ranked individuals. Furthermore, in line with the results of Study 1 people overweight private information when compared to social information derived from equally ranked persons.

Similar to Study 1 we compared the actual to the predicted confidence judgments of participants (see Figure 1B). Again, this test of the social influence model is performed purely on the predicted subjective probabilities that were derived from the model, which were estimated on the basis of participants' decisions. Thus, participants' confidence judgments were not used at all to fit the model. Again, the social influence model was able to predict people's confidence judgments very accurately.

Discussion of Study 2

The results of Study 2 support the view that individuals are affected by informational and authority influences. Consistently, the majority of participants made decisions that can be regarded as rational when considering the sequential decision problem from a Bayesian perspective. This holds for scenarios in which, according to a Bayesian analysis, the posterior probability of one disease is above .50. Authority influences could be observed when the decision of the medical director contradicted participants’ private information (compared to the baseline condition), independently of the corresponding posterior probability of the scenarios. The average proportion of decisions according to private information and probability judgments was consistently lower, illustrating the authority influences. With regard to the impact of authority influences supporting participants’ private information, only the analysis of participants’ probability judgments reveals a consistent pattern; that is, higher confidence in one’s own decision when the previous decision of the medical director is in line with participants’ private information. Finally, the results of our social influence model reveal that people treat public information differently due to its normative quality and independent of its validity. Moreover, the social influence model was also able to predict people's confidence judgments quite accurately, importantly without making use of the confidence data to estimate the model's parameter.

General Discussion

The primary goal of our studies was to examine how individuals' decisions are influenced by the decisions of others. Therefore, we tried to manipulate informational and authority influences by embedding a social decision task into different contexts. Using the cascade paradigm, we were able to trace back the effects of the two influence types on people's decisions. Study 1 shows that individuals do integrate socially inferred information to make a decision consistent with a Bayesian analysis. Study 2 shows the impact of authority and informational social influences on individual decision making. Authority influence affects people's judgment most when the decision of a higher ranked individual is opposed to participants' private information. In these types of situations, people show stronger conformity behavior and lower confidence in their own private information compared to situations in which they are confronted with opposing decisions of similar hierarchically ranked individuals. Additionally, we found consistent authority influences on participants' probability judgments when previous authority decisions supported participants' private information.

As a consequence, one can assume that the impact of authority influence should foster the emergence of information cascades. In Study 1 the majority of our participants decided in indifference situations according to their private information (on average 79.9% of all participants, Scenarios 10–12, see Table 2). In Study 2 the majority of our participants decided in indifference situations *against* their private information (on average 61.5% of all participants, Scenarios 36–40, see Table 6) when authority influences were exerted. Given the risk that two decision makers may have unfortunately obtained private information indicating the wrong state of affairs and that subsequent decision makers have followed them, the results of Study 2 reveal that only one authority decision will suffice to start a cascade, independently of subsequent privately obtained information.

The results of Studies 1 and 2 show that people apparently use social information to make decisions in a way that is generally consistent with a Bayesian perspective of the sequential decision problem. However, quantifying social influences with our computational model based on participants' choices shows that the weight people give to social and private information is context-dependent and therefore deviate from the pure Bayesian analysis that weights both kind of information equally. In line with recent studies on cascade behavior (e.g., Bernardo & Welch, 2001; Goeree et al., 2007; Nöth & Weber, 2003; Weizsäcker, 2010) we found that participants assigned higher weights to private information relative to public information within an *urn-and-balls*-setting (Study 1). Contrary to recent studies on cascade behavior, embedding the decision task into a real-life context reveals that people treat public information derived from higher ranked individuals more seriously than public information derived from equally ranked individuals whereas they overweight private information as compared to social information derived from lower ranked persons. Therefore, we argue that normative social influence cannot be neglected when analyzing the occurrence of information cascades in real-life settings. Moreover, the model, which was estimated only on the basis of participants' choices, was also able to predict people's confidence judgments. For Study 1 and Study 2 the model was able to explain why people's confidence judgments deviate from the posterior probabilities of the Bayesian account.

The current research throws new light on the motivational grounds of conformity by clarifying the role of informational social influence in relation to authority influence. The findings of both studies highlight the cognitive aggregation of available public and private information as a decisive factor in the occurrence of conformity. According to the informational hypothesis, people evaluate the validity of socially inferred information and integrate it to make a decision. Consistently, one can assume that people are principally influenced by information of others and that authority influences only marginally account for conformity behavior. However, in both studies we used a task in which participants' decisions could be objectively evaluated. Thus the impact of authority influence should have affected people's decisions less compared to tasks where objectively correct solutions are barely identifiable, if at all (e.g., judging the attractiveness of persons, Klucharev, Hytönen, Rijpkema, Smidts, & Fernández, 2009). Therefore, the results refer to social influence situations where the intellective properties of a task are salient (Kaplan, 1987; Kaplan & Miller, 1983; Laughlin, 1980).

From an applied perspective, our results reveal that the emergence of informational cascades can be fostered by authority influences. In particular, in situations in which the decisions of higher ranked individuals should have been given the same importance as those of other individuals due to equal decision accuracy, our results reveal that people still assign more importance to the decisions of the higher ranked individual. Here, the majority of our participants decided against their private information and thereby will start a cascade independently of subsequent privately obtained information. From this one may conclude that even when people act rationally according to a Bayesian perspective, the group of decision makers might not make good decisions as a whole. Thus, interventions to support sequential decision-making processes should focus more on changing the design of redundant systems rather than on changing the individual. Here it is important to change the structure of how individuals make decisions. For instance, one can think of systems where individuals first decide without knowing the decisions of their predecessors, and thereafter the single decisions are aggregated in a group context. This has the advantage that all available private information is integrated in the decision of the group.

Improving the reliability of sequential decision-making structures should also include reflections on the incentives that individuals expect. Our studies focused on situations in which people wanted to maximize their individual outcomes; however, social influence situations may differ with respect to their underlying incentive structure. On the one hand, there can be incentives for following the group regardless of being correct. On the other hand, social influence situations can provide incentives to follow the group and to make a correct decision. For example, Hung and Plott (2001) provided evidence on how information cascades developed when decision makers were positively rewarded when their personal decision was identical to the majority decision. They demonstrated that the attainment of a group goal led to a tendency to place more weight on public information than on private information. Therefore, it seems important to consider the incentives people expect in sequential decision-making structures and whether these goals correspond with their individual goals.

Our studies show that people cognitively integrate both private and public information for making decisions. They attach importance to the inferred information not solely based on their validity but also by taking into account the normative qualities of this information. Therefore, people make smart decisions that aim at being accurate and consistent with their social environment.

References

Allen, V. L. (1965). Situational factors in conformity. In L. Berkowitz (Ed.), *Advances in experimental social psychology* (Vol. 2, pp. 133-176). New York: Academic Press.

Allen, V. L., & Levine, J. M. (1971). Social support and conformity: The role of independent assessment of reality. *Journal of Experimental Social Psychology, 7*, 48-58.

Anderson, L. R., & Holt, C. A. (1997). Information cascades in the laboratory. *American Economic Review*, *87*, 847-862.

Asch, S. E. (1951). Effects of group pressure upon the modification and distortion of judgments. In H. Guetzkow (Ed.), *Groups, leadership, and men* (pp. 177-190). Pittsburgh, PA: Carnegie Press.

Asch, S. E. (1956). Studies of independence and conformity: 1. A minority of one against a unanimous majority. *Psychological Monographs*, *70*, 1-70.

Bernardo, A.E., & Welch, I. (2001). On the evolution of overconfidence and entrepreneurs. *Journal of Economics & Management Strategy, 10*, 301–330.

Bikhchandani, S., Hirshleifer, D., & Welch, I. (1992). A theory of fads, fashion, custom, and cultural change as informational cascades. *Journal of Political Economy*, *100*, 992-1026.

Bugental, D. B. (2000). Acquisition of the algorithms of social life: a domain-based approach. *Psychological Bulletin*, *126*(2), 187-219.

Chaiken, S., Wood, W., & Eagly, A. H. (1996). Principles of persuasion. In E. T. Higgins & A. W. Kruglanski (Eds.), *Social psychology: Handbook of basic principles* (pp. 702-742). New York: Guilford Press.

Cialdini, R. B., Bator, R. J., & Guadagno, R. E. (1999). Normative influences in organizations. In J. M. Levine, L. L. Thompson, & D. M. Messick (Eds.), *Shared cognition in organizations: The management of knowledge* (pp. 195-211). Hillsdale, NJ: Erlbaum.

Cialdini, R. B., & Goldstein, N. J. (2004). Social influence: Compliance and conformity. *Annual Review of Psychology*, *55*, 591-621.

Deutsch, M., & Gerard, H. B. (1955). A study of normative and informational influence upon individual judgment. *Journal of Abnormal and Social Psychology*, *51*, 629-636.

Edwards, W. (1968). Conservatism in human information processing. In B. Kleinmutz (Ed.), *Formal representation of human judgment* (pp. 17-52). New York: Wiley.

Epley, N., & Gilovich, T. (1999). Just going along: Nonconscious priming and conformity to social pressure. *Journal of Experimental Social Psychology*, *35*, 578-589.

Festinger, L. (1954). A theory of social comparison processes. *Human Relations*, *7*, 117-140.

Goeree, J. K., Palfrey, T. R., Rogers, B. W., & McKelvey, R. D. (2007). Self-correcting information cascades. *The Review of Economic Studies*, *74*(3), 733-762.

Grether, D. M. (1980). Bayes rule as a descriptive model: The representative heuristic. *Quarterly Journal of Economics*, *95*, 537-557.

Huck, S., & Oechssler, J. (2000). Informational cascades in the laboratory: Do they occur for the right reasons? *Journal of Economic Psychology*, *21*(6), 661-671.

Hung, A., & Plott, C. R. (2001). Information cascades: Replication and extension to majority rule and conformity-rewarding institutions. *American Economic Review*, *91*, 1508-1520.

Insko, C. A., Drenan, S., Solomon, M. R., Smith, P. B., & Wade, T. J. (1983). Conformity as a function of the consistency of positive self-evaluation with being liked and being right. *Journal of Experimental Social Psychology*, *19*, 341-358.

Kaplan, M. F. (1987). The influencing process in group decision making. In C. Hendrick (Ed.), *Group processes. Review of personality and social psychology* (Vol. 8). Beverly Hills, CA: SAGE.

Kaplan, M. F., & Miller, C. E. (1983). Group discussion and judgment. In P. B. Paulus (Ed.), *Basic group processes* (pp. 65-94). New York: Springer.

Kaplan, M. F., & Miller, C. E. (1987). Group decision making and normative versus informational influence: Effects of type of issue and assigned decision rules. *Journal of Personality and Social Psychology*, *53*, 306-313.

Klucharev, V., Hytönen, K., Rijpkema, M., Smidts, A., & Fernández, G. (2009). Reinforcement learning signal predicts social conformity. *Neuron*, *61*(1), 140-151.

Kruschke, J. K. (2010a). What to believe: Bayesian methods for data analysis. *Trends in Cognitive Sciences*, *14*(7), 293–300. doi:10.1016/j.tics.2010.05.001

Kruschke, J. K. (2010b). Bayesian data analysis. *Wiley Interdisciplinary Reviews: Cognitive Science*, *1*(5), 658–676. doi:10.1002/wcs.72

Kruschke, J. K. (2011). *Doing Bayesian data analysis: A tutorial with R and BUGS*. Burlington, MA: Academic Press.

Kübler, D., & Weizsäcker, G. (2005). Are longer cascades more stable? *Journal of the European Economic Association*, *3*, 330-339.

Laughlin, P. R. (1980). Social combination processes of cooperative problem-solving groups on verbal intellective tasks. In M. Fishbein (Ed.), *Progress in social psychology*. Hillsdale, NJ: Erlbaum.

Levine, J. M., & Russo, E. M. (1987). Majority and minority influence. In C. Hendrick (Ed.), *Review of personality and social psychology* (Vol. 8, pp. 13-53). Beverly Hills, CA: SAGE.

Milgram, S. (1974). Obedience to authority. *London: Tavistock (*1997).

Nöth, M., & Weber, M. (2003). Information aggregation with random ordering: Cascades and overconfidence. *The Economic Journal*, *113*(484), 166-189.

Pendry, L. F., & Carrick, R. (2001). Doing what the mob do: Priming effects upon conformity. *European Journal of Social Psychology, 31,* 83-92.

Phillips, L. D., & Edwards, W. (1966). Conservatism in a simple probability inference task. *Journal of Experimental Psychology*, *72*, 346-354.

Plummer, M. (2003). JAGS: A program for analysis of Bayesian graphical models using Gibbs sampling. In *Proceedings of the 3rd International Workshop on Distributed Statistical Computing (DSC 2003),* March 20-22, Vienna, Austria.

Sherif, M. (1935). A study of some social factors in perception. *Archives of Psychology*, *27*, 1-60.

Tajfel, H. (1969). Social and cultural factors in perception. In G. Lindzey & E. Aronson (Eds.), *The handbook of social psychology* (2nd ed., Vol. 3, pp. 315-394). Reading, MA: Addison-Wesley.

Weizsäcker, G. (2010). Do we follow others when we should? A simple test of rational expectations. *The American Economic Review, 100,* 2340-2360.

Wood, W. (2000). Attitude change: Persuasion and social influence. *Annual Review of Psychology*, *51*, 539-570.

Appendix: The Bayesian analysis of the sequential decision problem

According to a Bayesian analysis the posterior probability of Urn A being selected is determined by applying Bayes's theorem:

$$p(\mathrm{A} \mid n_{\mathrm{a}}, n_{b}) = \frac{p(n_{\mathrm{a}}, n_{b} \mid \mathrm{A})\, p(\mathrm{A})}{p(n_{\mathrm{a}}, n_{\mathrm{b}} \mid \mathrm{A})\, p(\mathrm{A}) + p(n_{\mathrm{a}}, n_{\mathrm{b}} \mid \mathrm{B})\, p(\mathrm{B})}$$

(A1)

where $p(n_{\mathrm{a}}, n_{\mathrm{b}} \mid \mathrm{A})$ is the likelihood of obtaining the number n_a and n_b of "a" and "b" signals given that Urn A was selected, where "a" speaks for Urn A and "b" speaks for Urn B. Analogously, the posterior probability of Urn B being selected given the number n_a and n_b of "a" and "b" signals can be determined, so that the ratio of the two posterior probabilities is defined as

$$\frac{p(\mathrm{A} \mid n_{\mathrm{a}}, n_{\mathrm{b}})}{p(\mathrm{B} \mid n_{\mathrm{a}}, n_{\mathrm{b}})} = \frac{p(n_{\mathrm{a}}, n_{\mathrm{b}} \mid \mathrm{A})}{p(n_{\mathrm{a}}, n_{\mathrm{b}} \mid \mathrm{B})} \cdot \frac{p(\mathrm{A})}{p(\mathrm{B})}$$

(A2)

Assuming equal a priori probabilities of the two urns being selected and taking the logarithm on both sides provides

$$\ln \frac{p(\mathrm{A} \mid n_{\mathrm{a}}, n_{\mathrm{a}})}{p(\mathrm{B} \mid n_{\mathrm{a}}, n_{\mathrm{b}})} = \ln \frac{p(n_{\mathrm{a}}, n_{\mathrm{b}} \mid \mathrm{A})}{p(n_{\mathrm{a}}, n_{\mathrm{b}} \mid \mathrm{B})},$$

(A3)

which can be rewritten as

$$\ln \frac{p(\mathrm{A} \mid n_{\mathrm{a}}, n_{\mathrm{b}})}{p(\mathrm{B} \mid n_{\mathrm{a}}, n_{\mathrm{b}})} = \sum_{i=1}^{n_{\mathrm{a}}} \ln \frac{p(\mathrm{a} \mid \mathrm{A})}{p(\mathrm{a} \mid \mathrm{B})} + \sum_{i=1}^{n_{\mathrm{b}}} \ln \frac{p(\mathrm{b} \mid \mathrm{A})}{p(\mathrm{b} \mid \mathrm{B})}.$$

(A4)

Footnotes

1. These decisions are also in line with a Bayesian analysis, if one takes into account that people occasionally make mistakes. If one assumes that with a small probability the preceding persons in the sequential decision situation might have chosen the wrong urn, this increases the posterior probability above .50 in favor of one's own information, so that a decision in line with one's own signal should be made. For instance, we assumed that two initial A decisions imply two "a" draws, which means that if the second decision maker had a private "b" signal we assume that she would have selected Urn B. In the case of a "b" signal for the second decision maker the posterior probabilities are equal for both urns, but because the second decision maker cannot be absolutely sure that the first decision maker has not mistakenly selected Urn A, it is reasonable for her to go with her own signal. Nevertheless, even if one assumes that the second decision maker decided randomly when she got a "b" signal, the third decision maker should still predict Urn A (against his own "b" signal), as the posterior probability for Urn A is $p = .56$.

2. The 95% highest density interval (HDI) is a way to summarize posterior distributions used in Bayesian statistics. According to Kruschke (2011) the 95% HDI can be defined as "[…] an interval that spans 95% of the distribution, such that every point inside the interval has higher believability than any point outside the interval" (p. 85).

3. To guarantee that the three importance parameters sum up to 3 we applied the following procedure: In a first step two independent values (β_1 and β_2) were sampled from uniform beta distributions ($\alpha = 1$, $\beta = 1$). The smaller value was then multiplied by 3 and defines β_(public equal rank). To get β_(public higher rank) we subtracted the smaller value from the higher value and multiplied the result by 3. Finally, the importance parameter for the private information was derived by subtracting β_(public equal rank) and β_(public higher rank) from 3.

Tables

Table 1

Participants' decisions and probability judgments for the nine decision scenarios of Study 1 in which, according to a Bayesian solution, the posterior probability of one urn being chosen is above .50.

Scenario	Previous decisions	Private information favors	Posterior probability	Choices for the most likely urn (%)[a]	Average probability judgment
1	Urn A	Urn A	.80 for A	91.3	.66
2	Urn A; Urn A	Urn A	.89 for A	90.0	.74
3	Urn A; Urn B	Urn A	.67 for A	91.1	.62
4	Urn A; Urn A; Urn A	Urn A	.89 for A	85.0	.75
5	Urn A; Urn B; Urn A	Urn A	.80 for A	85.0	.69
6	Urn A; Urn A	Urn B	.67 for A	71.3	.54
7	Urn A; Urn B	Urn B	.67 for B	95.0	.66
8	Urn A; Urn A; Urn A	Urn B	.67 for A	79.7	.65
9	Urn A; Urn B; Urn B	Urn B	.80 for B	93.8	.74

[a] Most likely according to a Bayesian analysis

Table 2

Participants' decisions and probability judgments for the three decision scenarios in Study 1 in which a Bayesian analysis leads to an indifference situation; that is, the posterior probability for both urns being .50.

Scenario	Previous decisions	Private information favors	Posterior probability	Choices for the urn favored by private signal (%)	Average probability judgment
10	Urn A	Urn B	.50	90.0	.60
11	Urn A; Urn B; Urn B	Urn A	.50	65.0	.60
12	Urn A; Urn B; Urn A	Urn B	.50	84.8	.58

Table 3

Participants' decisions and probability judgments for the 13 decision scenarios of Study 2, in which the posterior probability of one disease according to a Bayesian analysis is .67.

Scenario	Previous diagnosis	Private information favors	Posterior probability	Participants choosing the most likely disease (%)[a]	Participants' average probability judgment	Average proportion of decisions according to private information	Average proba-bility judgment
Baseline Scenarios (no previous decision of the MD)							
1	AP: A, AP: S	A	.67 for A	95.0	0.70		
2	AP: A; AP: S	S	.67 for S	92.5	0.66		
3	AP: A; AP: A	S	.67 for A	75.0	0.60	.59	.65
4	AP: A; AP: A; AP: A	S	.67 for A	75.0	0.66		
Scenarios where the decision of the MD favors participants' private information							
5	MD: A, AP: S	A	.67 for A	95.0	0.72		
6	AP: A; MD: S	S	.67 for S	92.5	0.68	.94	.69
Scenarios where the decision of the MD speaks against participants' private information							
7	MD: A; AP: S	S	.67 for S	87.5	0.61		
8	AP: A, MD: S	A	.67 for A	82.5	0.64		
9	MD: A; AP: A	S	.67 for A	82.5	0.67		
10	AP: A; MD: A	S	.67 for A	82.5	0.65	.36	.66
11	MD: A; AP: A; AP: A	S	.67 for A	87.5	0.70		
12	AP: A; MD: A; AP: A	S	.67 for A	87.5	0.72		
13	AP: A; AP: A; MD: A	S	.67 for A	85.0	0.71		

Note. AP = Assistant physician; MD = Medical director; A = Appendicitis; S = Sigma diverticulitis.

[a] Most likely according to the Bayesian analysis

Table 4

Participants' decisions and probability judgments for the 10 decision scenarios of Study 2 in which the posterior probability of one disease according to a Bayesian analysis is .80.

Scenario	Previous diagnosis	Private information favors	Posterior probability	Participants choosing the most likely disease (%)	Participants' average probability judgment	Average proportion of decisions according to private information	Average proba-bility judgment
Baseline Scenarios (no previous decision of the MD)							
14	AP: A	A	.80 for A	97.5	0.80		
15	AP: A, AP: S; AP: A	A	.80 for A	95.0	0.77	.97	.77
16	AP: A; AP: S; AP: S	S	.80 for S	97.5	0.75		
Scenarios where the decision of the MD favors participants' private information							
17	MD: A	A	.80 for A	97.5	0.83		
18	MD: A, AP: S; AP: A	A	.80 for A	95.0	0.79		
19	AP: A, AP: S; MD: A	A	.80 for A	97.5	0.79	.97	.79
20	AP: A; MD: S, AP: S	S	.80 for S	95.0	0.76		
21	AP: A; AP: S; MD: S	S	.80 for S	97.5	0.77		
Scenarios where the decision of the MD speaks against participants' private information							
22	AP: A, MD: S; AP: A	A	.80 for A	90.0	0.72		
23	MD: A; AP: S; AP: S	S	.80 for S	82.5	0.63	.86	.67

Note. AP = Assistant physician; MD = Medical director; A = Appendicitis; S = Sigma diverticulitis.

[a] Most likely according to the Bayesian analysis

Table 5

Participants' decisions and probability judgments for the seven decision scenarios of Study 2 in which the posterior probability of one disease according to a Bayesian analysis is .89.

Scenario	Previous diagnosis	Private information favors	Posterior probability	Participants choosing the most likely disease (%)a	Participants' average probability judgment	Average proportion of decisions according to private information	Average proba-bility judgment
Baseline Scenarios (no previous decision of the MD)							
24	AP: A; AP: A	A	.89 for A	100.0	0.84		
25	AP: A; AP: A; AP: A	A	.89 for A	97.5	0.86	.98	.85
Scenarios where the decision of the MD favors participants' private information							
26	MD: A; AP: A	A	.89 for A	100.0	0.87		
27	AP: A; MD: A	A	.89 for A	100.0	0.85		
28	MD: A; AP: A; AP: A	A	.89 for A	100.0	0.89	1	.88
29	AP: A, MD: A; AP: A	A	.89 for A	100.0	0.88		
30	AP: A, AP: A; MD: A	A	.89 for A	100.0	0.88		

Note. AP = Assistant physician; MD = Medical director; A = Appendicitis; S = Sigma diverticulitis.

[a] Most likely according to the Bayesian analysis

Table 6

Participants' decisions and probability judgments for the decision scenarios in Study 2, where the posterior probabilities of the medical diagnosis task predict an indifference situation between the available options.

Scenario	Previous diagnosis	Private information favors	Participants' diagnosis according to their private information (%)	Participants' average probability judgment	Average proportion of decisions according to private information	Average probability judgment
Baseline Scenarios (no previous decision of the MD)						
31	AP: A	S	70.0	0.62		
32	AP: A, AP: S; AP: S	A	40.0	0.63	.57	.63
33	AP: A; AP: S; AP: A	S	60.0	0.66		
Scenarios where the decision of the MD favors participants' private information						
34	MD: A, AP: S; AP: S	A	75.0	0.68		
35	AP: A; MD: S, AP: A	S	60.0	0.69	.67	.69
Scenarios where the decision of the MD speaks against participants' private information						
36	MD: A	S	37.5	0.62		
37	AP: A, MD: S; AP: S	A	40.0	0.63		
38	AP: A, AP: S; MD: S	A	30.0	0.66	.39	.65
39	MD: A; AP: S; AP: A	S	47.5	0.66		
40	AP: A; AP: S; MD: A	S	37.5	0.68		

Note. AP = Assistant physician; MD = Medical director; A = Appendicitis; S = Sigma diverticulitis.

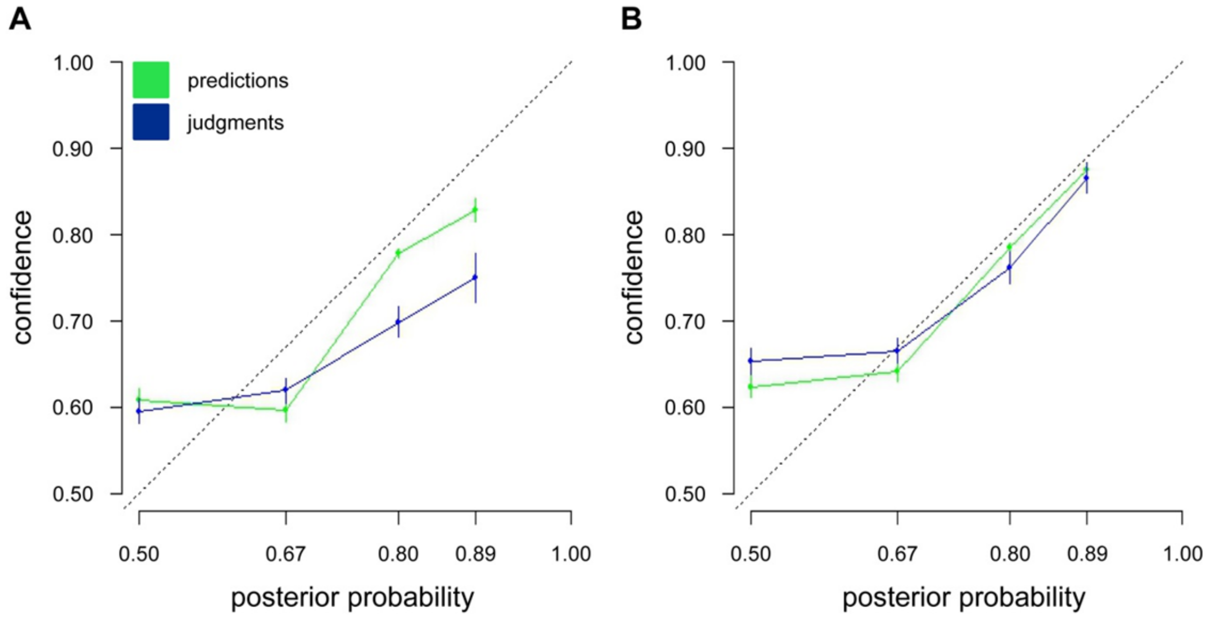

Figure 1. Empirically observed versus predicted confidence judgments for Study 1 (A) and Study 2 (B). The general pattern of confidence judgments (blue) is accurately captured by the predictions of the model solely derived from participants' choices (green) for Study 1 (A) and Study 2 (B).

Note. Confidence judgments on the dashed line are in accordance with the Bayesian solution.