

DISSERTATION

zur Erlangung des akademischen Grades
Doktor der Naturwissenschaften (Dr. rer. nat.)

How social learning strategies boost or undermine decision making in groups

eingereicht am Fachbereich Erziehungswissenschaft und
Psychologie der Freien Universität Berlin
von **Alan Novaes Tump, M.Sc.**

Berlin, 2019

Promotionskommission

Erstgutachter: Prof. Dr. Ralph Hertwig

Zweitgutachter: Prof. Dr. Hauke Heekeren

Prof. Dr. Jens Krause

Prof Dr. Dirk Ostwald

Dr. Julia Rodríguez Buritica

Tag der Disputation: 06. Juli 2020

Summary

Social interactions resulting in emergent collective behaviour play a key role in almost all layers of society, from local, small-scale interactions, such as people crossing the street, to global, large-scale interactions, such as the spread of fake news on online platforms. In our digital and interconnected world, it is increasingly important to understand the emergence of beneficial or detrimental collective dynamics. The characteristics of such dynamics are expected to depend greatly on the nature of information individuals have personally acquired and how they learn from others. Yet, how the decision-making processes shape the resulting collective dynamics remains poorly understood. When do individuals seek more information from social sources? How do individuals reap the benefits when navigating in social environments, and when do they fail to do so? This dissertation aims to answer these questions extending established theories and frameworks from individual decision-making into the social realm. This approach allows for the operationalization of personal and social information in a theory-driven manner, thereby achieving a deeper understanding of the individual-level decision process.

The first chapter provides a introductory overview of the interplay between personal information use, social learning strategies and collective dynamics, and introduces the key theories and models I will expand on in this dissertation. In Chapter 2, inspired by Brunswick's lens model, I investigate how individuals form beliefs about the meaning of ecological structures (i.e., cues). Here, participants had to categorize images based on multiple cues, the meaning of which had to be learned over trials. I showed that participants observing the same cues formed different beliefs about the cue meanings. This diversity in cue beliefs is, in turn, an important process governing the quality of social information. The greater this diversity, the more independent personal information is, and the stronger the potential for social information use. Participants, however, failed to realize the full potential of this diversity because they only changed their personal decisions if a large majority disagreed with them. Simulating different strategies of social information use, I show that this reliance on strongly agreeing majorities impedes individuals from benefiting from diversity. This chapter thus identifies diversity in cue beliefs as an important factor allowing individuals in groups to benefit from the wisdom of each other, while simultaneously highlighting the importance of the individuals' social learning strategies to exploit this diversity.

Chapter 3 dives deeper into the social learning strategies individuals use. By carefully controlling the social information displayed to participants, the study in this chapter provides an in-depth analysis of social learning strategies. Participants were confronted with an estimation task. They first provided an independent estimate, after which they observed estimates of others. Using Bayesian modelling techniques, I show that the incorporation of others' opinions strongly depends on how consistent those opinions are with an individual's own opinion and the degree of agreement among others. Individuals also strongly differ in the social learning strategies they use. These results elucidate what aspects are conducive for people to change their minds and contribute to the understanding of how individuals' social information use shapes opinion and information dynamics in our interconnected society.

In Chapter 4, I embed individuals in a temporal dynamic system which allows the investigation of the use of information in interaction with the emergent collective dynamic. Here, my focus is on social interactions where multiple people make decisions sequentially and thereby are simultaneously emitters and receivers of social information. To shed light on the unfolding dynamic in such settings, I will introduce the social drift-diffusion model (DDM). The model allows the investigation of the cognitive processes underlying the integration of personal and social information dynamically over time, and the subsequent collective dynamic. Analysis of the data shows that correct information spreads when the participants' confidence reflects accuracy and when more confident participants decide faster. Under these conditions, later-deciding participants are likely to adopt social information and thereby to amplify the correct signal provided by early-deciding participants. The social DDM successfully captures all the key dynamics observed in the social system, revealing the cognitive underpinnings of information cascades in social systems.

The general principles of personal and social information use that emerge from Chapter 4 allow to investigate the optimal behaviour when deciding sequentially. In Chapter 5, I develop an agent-based version of the social DDM and embed it in evolutionary algorithms, allowing the identification of evolutionarily advantageous strategies. I show that the individuals' decision time should depend on the quality of information, with the most accurate individuals deciding first. For all later-deciding individuals it is evolutionary advantageous to imitate the (often accurate) first decision. When introducing asymmetric error costs, single individuals should develop response biases to avoid the more costly error. In groups, however, such response biases can have dramatic consequences, as these biases are likely to be amplified in the group. As a result, individuals in large groups should use much weaker response biases to benefit from social information. I conclude that individuals facing asymmetric error costs in social environments need to carefully trade off the expressed response bias and sensitivity to social information to avoid the more costly error but simultaneously benefit from the collective.

Overall, this thesis deepens our understanding of social dynamics by accounting for individual-level decision-making processes across various choice problems. I show that the behaviour of individuals in social environments can significantly differ depending on the personal information individuals possess and the strategies individuals use. Furthermore, I highlight the importance of accounting for such differences to predict the emergence of beneficial or detrimental dynamics in social environments.

Zusammenfassung

Soziale Interaktionen spielen auf fast allen Ebenen der Gesellschaft eine zentrale Rolle, von lokalen Interaktionen, wie dem Überqueren der Straße, bis hin zu globalen Interaktionen, wie der Verbreitung von falschen Nachrichten auf Online-Plattformen. In unserer digitalen und eng vernetzten Welt wird es immer wichtiger, die Entstehung nützlicher oder schädlicher kollektiver Dynamiken zu verstehen und vorherzusagen. Die Eigenschaften der kollektiven Dynamiken hängen stark von der Art der Informationen ab, die die einzelnen Personen besitzen, und wie diese Personen voneinander. Wie die individuellen Entscheidungsprozesse die daraus entstehende kollektive Dynamik prägen, bleibt jedoch wenig erforscht. Wann suchen Personen gezielt Informationen aus ihrem sozialen Umfeld? Wann profitiert der Einzelne davon, und wann nicht? Die vorliegende Arbeit will diese Fragen mit Hilfe etablierter Theorien zur Entscheidungsfindung, die auf den sozialen Kontext erweitert werden, beantworten. Dieser theoretisch fundierte Ansatz ermöglicht ein tieferes Verständnis von Entscheidungsprozessen in Gruppen.

Das erste Kapitel bietet einen theoretischen Überblick über das Zusammenspiel von persönlichen Informationen, Strategien zum sozialen Lernen und kollektiver Dynamik. Außerdem stellt es die wichtigsten Theorien und Modelle vor, die in dieser Dissertation erweitern werden. Inspiriert durch Brunswicks Linsenmodell wird in Kapitel 2 untersucht wie Personen die Bedeutung von Hinweisen bewerten. Ich zeige, dass obwohl Personen dieselben Situationen beobachten, sie unterschiedliche Bewertungsansichten dieser Hinweise entwickeln. Die Vielfalt der Ansichten ist wiederum eine wichtige Voraussetzung für qualitativ gute soziale Informationen. Je höher die Ansichtsvielfalt, desto unabhängiger sind die Personen und desto größer ist das Potenzial für eine erfolgreiche Nutzung sozialer Informationen. Die Probanden in dem Versuch konnten jedoch nicht das volle Potenzial dieser Vielfalt ausschöpfen, da sie sich nur auf große Mehrheiten verließen. Probanden änderten ihre persönliche Entscheidung nur, wenn eine große Mehrheit anderer Probanden nicht mit ihrer Meinung übereinstimmte. Durch die Simulation verschiedener Strategien zeige ich, dass Personen mehr von der Ansichtsvielfalt profitiert hätten, wenn sie sich auch auf kleinere Mehrheiten verlassen hätten. Dieses Kapitel zeigt daher, dass die Bewertungsvielfalt von Hinweisen erlaubt, dass Personen von der Weisheit der Vielen profitieren können. Es betont aber gleichzeitig wie wichtig es ist die Lernstrategien der Individuen zu beachten um zu Verstehen wann diese Vielfalt nützlich ist.

Kapitel 3 geht tiefer auf die sozialen Lernstrategien ein, die Individuen anwenden. Die genaue Kontrolle der sozialen Informationen, die den Teilnehmern im Versuch gezeigt werden, ermöglicht in diesem Kapitel eine detaillierte Analyse der sozialen Lernstrategien. Die Teilnehmer der Studie schätzten zunächst selbst die Anzahl von Objekten in einem Bild, woraufhin sie Schätzungen anderer beobachteten. Mit Hilfe bayesianischer Statistik zeige ich, dass die Einbeziehung der Meinungen anderer in hohem Maß davon abhängt, wie sehr diese die eigene Meinung wieder spiegeln. Außerdem werden andere Einschätzungen stärker gewertet, wenn diese mehr miteinander Übereinstimmen. Personen unterscheiden sich sehr in den Strategien, mit denen sie von anderen

lernen. Die Ergebnisse zeigen auf unter welchen Umständen Personen ihre Meinung ändern, und tragen zum Verständnis bei, wie die Nutzung sozialer Informationen Meinungsdynamiken in unserer Gesellschaft formt.

Als nächsten Schritt biete ich Individuen in ein zeitliches dynamisches System ein, um die Interaktion zwischen Informationsnutzung und den entstehenden kollektiven Dynamiken zu untersuchen. Hier konzentriere ich mich auf soziale Interaktionen, bei denen mehrere Personen hintereinander entscheiden und damit gleichzeitig Sender und Empfänger von sozialen Informationen sind. Um die sich entwickelnde Dynamik in solchen Gruppen zu verstehen, stelle ich das *social drift-diffusion Model* (social DDM) vor. Das Modell ermöglicht die Untersuchung der zugrunde liegenden kognitiven Prozesse bei der dynamischen Informationsaufnahme über Zeit sowie der entstehenden kollektiven Dynamik. Empirisch stelle ich heraus, dass sich korrekte Informationen verbreiten, wenn die Zuversicht auf die Richtigkeit der eigenen Informationen diese auch wirklich widerspiegelt. Zusätzlich sollten die zuversichtlichen Personen sich schneller entscheiden. Unter diesen Bedingungen ist es wahrscheinlich, dass später entscheidende Personen soziale Informationen übernehmen und damit das richtige Signal früh entscheidender Personen weiter verstärken. Das social DDM erfasst erfolgreich alle beobachteten Schlüsseldynamiken und gibt Einblicke in kognitive Prozesse bei Informationskaskaden.

Die beobachteten allgemeinen kognitiven Prozesse ermöglichen es zu untersuchen, wie Personen sich optimal Verhalten sollten. In Kapitel 5 biete ich das etablierte social DDM in evolutionäre Algorithmen ein, die es ermöglichen, evolutionär vorteilhafte Strategien zu identifizieren. Ich zeige, dass der Zeitpunkt der Entscheidung von der Qualität der Informationen abhängen sollte, wobei die Person mit den meisten Informationen zuerst entscheiden sollte. Für alle späteren Entscheider ist es von Vorteil, die (oft richtige) erste Entscheidung nachzuahmen. Bei der Einführung asymmetrischer Fehlerkosten sollten Personen eine Antworttendenz entwickeln, um den kostspieligeren Fehler zu vermeiden. In Gruppen können solche systematischen Antworttendenzen jedoch dramatische Folgen haben, da sich diese in der Gruppe wahrscheinlich verstärken werden. Infolgedessen sollten Personen in großen Gruppen deutlich schwächere Tendenzen aufweisen, um von sozialen Informationen zu profitieren. Ich komme zu dem Schluss, dass Personen, die mit asymmetrischen Fehlerkosten konfrontiert sind, die Antworttendenzen und die Sensibilität für soziale Informationen sorgfältig abwägen müssen, um die kostspieligeren Fehler zu vermeiden und gleichzeitig vom Kollektiv zu profitieren.

Insgesamt vertieft diese Arbeit unser Verständnis von sozialen Dynamiken, indem sie Entscheidungsprozesse auf individueller Ebene über viele verschiedene Entscheidungsprobleme hinweg untersucht. Ich zeige, dass das Verhalten von Individuen im sozialen Umfeld je nach den persönlichen Informationen, die sie besitzen, und den Strategien, die sie anwenden, erheblich variieren kann. Darüber hinaus betone ich, wie wichtig es ist, solche Unterschiede zu berücksichtigen, um das Entstehen einer positiven oder negativen Dynamik im sozialen Umfeld vorherzusagen.

Acknowledgements

First of all, this thesis would not have been possible without Ralf Kurvers. Ralf, I am a very lucky student to have you as my supervisor and mentor. I cannot emphasize enough how grateful I am that you patiently gave me the freedom to find and pursue my interests while still nudging me into the right directions. You have been an inexhaustible source of support and advice.

I am grateful to Ralph Hertwig for giving me the opportunity to do my PhD at the Center for Adaptive Rationality (ARC). Thank you very much for challenging and supporting me throughout my time at ARC. I am very grateful that I had the opportunity to work and be inspired in such an interdisciplinary group.

Special thanks go to Tim Pleskac for the many discussions during which I have learned a lot, in particular concerning computational modelling. I want to thank Lucas Molleman and Wouter van den Bos for all the fruitful discussions and the opportunity to visit you and present my work in Amsterdam. I also want to thank Max Wolf for our encouraging and inspiring meetings. I am very grateful to have the possibility to work with such enthusiastic researchers.

Many thanks must go to the ARC Assistants, Katja Münz, Petra Siemers-Hering and Maren Kutscha for making sure that everything in ARC runs as smoothly as possible and thereby foster the inspirational work atmosphere at ARC.

Special thanks go to Kyanoush, Lou, Ralf, Simon, Stefan, Charley, Bertrand, Christina, Nika, Shahar, Phillipp, Juan, Ruben and everyone else at ARC for all the great time at the institute as well as during the retreats. I would also like to thank Jens Krause, Stefan Krause, Lysanne Snijders and everyone who joined the yearly research expeditions to Trinidad for this enriching experience.

Finally, I would like to thank my family and friends. Many thanks go to my Crew Till, Jan Philipp and Patrick for all the backup and the great time we share. I want to take this opportunity to thank my parents Ilma and Rainer and my brother Lukas for their constant support. I am very proud to have such an awesome family that always has my back. Finally, I want to thank my wife, Marie, for being my companion. Marie, your understanding, humor and encouragement have been the true source of inspiration through all those years.

Contents

1	General Introduction	3
2	Individuals fail to reap the collective benefits of diversity because of over-reliance on personal information	18
3	Strategies for integrating disparate social information	19
4	Wise or mad crowds? The cognitive mechanisms underlying information cascades	39
5	Adaptive decision rules in groups under asymmetrical error costs	64
6	Synthesis and Future Directions	82
	Appendices	93
	B Supplementary Material to Chapter 3	94
	C Supplementary Material to Chapter 4	112
	D Supplementary Material to Chapter 5	123
	Declaration of Independent Work	125

Chapter 1

General Introduction

Social learning

Throughout their lives, humans and animals, face problems that involve uncertainty. Individuals thereby often have incomplete information but are aware that others might have complementary information. Hence, in many situations, observing the behaviour or decisions of others to acquire, so-called, social information is a beneficial strategy (Galef, 2009; Danchin et al., 2004). What to eat, which books to buy, whether to cross the street, or whether to get vaccinated, all of these choices are likely to be influenced by the behaviour of others. The predisposition to learn from the behaviour of others is a fundamental element of cognition and has been investigated by social psychologists and animal behaviourists for centuries. Social learning is assumed to be advantageous because individuals can save time and effort on gathering personal information. Consumers do not have to evaluate all potential products or animals all potential foods. Instead, they could simply copy the behaviour of other individuals.¹

Despite the intuitive appeal of this premise, copying others indiscriminately is considered a poor strategy for individuals (Laland, 2004), and can even have negative consequences for the population or society as a whole (Giraldeau et al., 2002). The inefficiency of indiscriminatory copying lies in the parasitic nature of this strategy. As pointed out by Rogers (1988), a population relying exclusively on copying would lack individuals providing reliable information, a finding commonly known as Rogers' paradox (Rogers, 1988; Boyd and Richerson, 1995). In other words, if everyone just follows, nobody gets anywhere. Rogers' paradox inevitably raises the question: How do individuals reap the benefits of social environments? Laland (2004) answers this question by arguing that individuals need to acquire (social) information strategically. He introduces a collection of adaptive heuristics or *social learning strategies* which individuals should adaptively rely on, dividing them into two groups.

The first group answers the question of **when** to use social information, focusing on an individual's personal information. Individuals should rely on social information if they face high uncertainty and if gathering personal information is costly (Boyd and Richerson, 1996, 1985). A large body of literature has indeed shown that individuals strongly differ in the propensities for social learning. Individuals who are inexperienced (Baude et al., 2008; Kawaguchi et al., 2007; Morgan et al., 2012; Smolla et al., 2016) or lack confidence (Morgan et al., 2012; Toelch et al., 2014) rely more on social information than others. Also, characteristics of the task itself such

¹Copying here refers to any kind of social learning.

as increasing task difficulty (McElreath et al., 2005) or cost of personal information acquisition (Webster and Laland, 2008; Moussaïd et al., 2016) have been shown to increase the tendency to learn from others. **When** individuals use social information crucially influences the quality of information within a social environment. Strong reliance on personal information might hinder the spread of high-quality social information, while strong reliance on social information might hinder the acquisition of personal information (Laland, 2004).

The other group of strategies targets the immediate social environment, focusing on the question: From **whom** should we learn? Obviously, learning from informed others is likely to be a very effective strategy. Yet, it is often difficult to deduce the quality of information on which observed decisions are based. Instead, the quality of social information needs to be inferred by observing indirect cues of success. Previous research has shown that individuals, indeed, rely on such cues, including past performance, expressed confidence, wealth, health, age, and hierarchical status (Laland, 2004; Drea and Wallen, 1999; Schlag, 1998; Moussaïd et al., 2013; Deaner et al., 2005; Yaniv and Kleinberger, 2000; Reeb, 2001).

Horner et al. (2010), for example, found that chimpanzees (*Pan troglodytes*) preferably learn from older and higher ranking individuals—characteristics which covary with success. Such strategies that involve using cues to identify accurate individuals often performs well. However, the downside is that these strategies fail as soon as the cues supposedly indicating accuracy are unreliable (Herzog et al., 2019).

A more robust strategy is to copy the behaviour expressed by the majority (Laland, 2004; Hastie and Kameda, 2005; Van Vugt, 2006), also known as conformist behaviour (Morgan and Laland, 2012; Henrich and Boyd, 1998; Bond and Smith, 1996; Morgan et al., 2012). One of the earliest demonstrations of its power was provided by the mathematician and philosopher Nicolas de Condorcet. In a nutshell, his Jury Theorem states that the accuracy of a majority of voters increases with the number of voters (i.e., individuals making choices), provided that these voters perform better than chance (Condorcet, 1785), an effect also known as the “Wisdom of Crowds” (Surowiecki, 2004; Laan et al., 2017; Bang and Frith, 2017; Mannes et al., 2014; Herzog et al., 2019). Past studies have found evidence for conformist behaviour in a wide range of animal taxa including insects (Bridges and Chittka, 2019), birds (Aplin et al., 2015), fish (Day et al., 2001; Pike and Laland, 2010; Kendal et al., 2004), rats (Konopasky and Telegdy, 1977; Galef and Whiskin, 2008), monkeys (Dindo et al., 2009), and great apes (Whiten et al., 2005). Conformity has also been studied extensively in humans (Bond and Smith, 1996; McElreath et al., 2005; Toelch and

Dolan, 2015; Moussaïd et al., 2016; Morgan et al., 2012; Milgram et al., 1969).

Arguably, the most famous experiments on conformity were conducted by Solomon Asch in the 1950s (Asch, 1956; Asch and Guetzkow, 1951). Here, participants performed a simple perceptual judgment task in the presence of confederates, who were instructed to often make incorrect judgments. Asch found that participants regularly abandoned their own judgment when facing a majority of confederates disagreeing with them. The study thereby provided an intriguing example of the persuasive power of (even wrong) majorities.² Yet, it also reveals the detrimental effect of relying on majorities when these are likely to be wrong. In highly spatially and temporally fluctuating environments, for example, the behaviour of the population might not be adapted to the current conditions and, hence, relying on the majority would be a bad choice (Henrich and Boyd, 1998; Feldman et al., 1996). Another important requirement for the success of this strategy is diversity (Kaniovski, 2010; Ladha, 1992, 1995; Sekiguchi and Ohtsuki, 2015). If individuals have access to diverse information (Sorkin et al., 2001), or use diverse inference strategies (Fujisaki et al., 2018), they are expected to make diverse errors and thereby promote the Wisdom of Crowds effect. In this thesis I will investigate a further mechanism promoting diversity and show which learning strategies allow individuals to make use of it.

While past research has identified various factors influencing social learning and proposed various verbal concepts explaining these findings, it has been challenged to find a single holistic, theoretically-grounded model. Arguably, the most successful attempt is Latané's (1981) social impact theory. He operationalizes social influence by describing it as a social force—analogous to a physical force such as light, gravity or magnetism—that acts within its force field (i.e., social surrounding). Latané identified three key factors that determine the magnitude of social influence:

Strength: determined by the characteristics of the social source (e.g., past accuracy or age).

Immediacy: determined by the proximity of the social source in time and space.

Number of sources: the number of other individuals present.

Having a formal model—with principles borrowed from physics—allowed social impact theory to accurately predict the diminishing additional effect of increasing the number of social sources as observed in Asch's (1951) and Milgram's (1969) experiments. The variables of the model (i.e.,

²Note that the participants' behaviour in Asch's experiments can also be explained by a motivation to conform to the group's norm rather than informational gain (Deutsch and Gerard, 1955).

strength, immediacy and number) thereby account for strategies targeting the social environment (i.e, who to copy). However, they are blind to the information individuals possess (i.e, when to copy) and by that miss a crucial aspect of social learning strategies. This thesis will provide new theories that formalize both: when and from whom to learn.

Collective dynamics

The importance of social influence goes beyond an individual's personal experience. Judgements or behaviours can spread through a population via social influence, thereby affecting the dynamics of collective behaviour. Social influence, for example, is considered to be a driver of self-organization (Deneubourg et al., 1990; Kurvers et al., 2015; Couzin and Krause, 2003), swarm intelligence (Krause et al., 2010; Couzin et al., 2005), and cultural evolution (Danchin et al., 2004; Morgan et al., 2012; Kempe and Mesoudi, 2014). In addition, it is a key mechanism driving the dynamics in a range of societal areas, including fashion (Salganik et al., 2006), the spread of violence (Slutkin et al., 2015), political mobilisation (Battaglini, 2005; Bond et al., 2012), consumer preferences (Chen, 2008), and financial markets (Shiller, 2002; Welch, 2000).

Social influence can positively affect collective dynamics, but it can also threaten collective systems. On the one hand, collective dynamics, fueled by social influence, allows bees to choose the best nest site (List et al., 2009), termites to build their nest (Deneubourg, 1977), pigeons to find their way home (Biro et al., 2006), and modern societies to rise, via the diffusion of innovation (Goldstone and Gureckis, 2009). On the other hand, collective dynamics can lead to the spread of fake news (Xiong and Liu, 2014; Vosoughi et al., 2018), maladaptive herding (Toyokawa et al., 2019), echo chambers, and filter bubbles (Lewandowsky et al., 2017). In an increasingly digital and interconnected society with platforms such as Twitter, Facebook, Reddit or Amazon, which are designed for information to propagate, it is important to understand the mechanisms underlying such systems.

An influential approach to understanding the spread of information is the concept of information cascades (Bikhchandani et al., 1998; Banerjee, 1992; Anderson and Holt, 1997): A single individual observing a small group expressing a particular behaviour is likely to adopt this behaviour, especially in the absence of personal information. The next individual will now face an even larger group expressing a certain behaviour and is, therefore, even more likely to adopt this behaviour. In an extreme case, everyone adopts the behaviour originally expressed by only a small group. Understanding the dynamics of such cascades, and particularly when these cascades go

wrong, is important for a wide range of social systems.

Whether information cascades promote positive or negative outcomes strongly depends on the initial (early) behaviour and the social learning strategies individuals use. Previous research on cascades has suggested that both decision timing and social learning strategies should be based on the quality of personal information (Chamley and Gale, 1994; Gul and Lundholm, 1995; Zhang, 1997; Ziegelmeyer et al., 2005). Hence, to predict collective behaviour it is important to take personal information, social learning strategies and their interaction into account. Past researchers have developed a number of models to describe collective dynamics. Many target the question of how individuals coordinate spatially and, for example, investigate collective motion in fish (Couzin et al., 2005), birds (Flack et al., 2015), monkeys (Strandburg-Peshkin et al., 2013) or pedestrians (Helbing et al., 2002; Hoogendoorn and Daamen, 2007). Another class of—non-spatial—models target the spread of information conceptually, for example, in sequential decisions (Sumpter and Pratt, 2008; Deneubourg et al., 1990; Bikhchandani et al., 1998), innovation diffusion (Rogers, 2004), social networks (Abrahamson and Rosenkopf, 1997) or opinion dynamics (Lorenz, 2007; Granovetter, 1978).

These models have been very instrumental in increasing our understanding of emergent collective behaviour by investigating how simple principles of social interactions shape collective dynamics. Yet, they have received criticism from cognitive psychologists who pointed out that they often have a top-down or bird’s-eye perspective, where individuals are treated as units or atoms with simple interaction rules. In short, they focus on patterns, not on people (Raafat et al., 2009). While doing so, they regularly oversimplify the influence of the cognitive processes on the unfolding dynamics. Similarly, research on social learning strategies has identified choice rules individuals should rely on, but usually ignores the cognitive processes implementing them (Heyes, 2016a,b). In parallel, an increasing number of studies emphasize the ability of (individual) decision-making models to account for such cognitive processes under social influence (Germar et al., 2014; Toelch and Dolan, 2015; Toelch et al., 2018; Bang and Frith, 2017). Researchers have, thus, called to for an integrative approach to link the different levels of decision-making in groups (Raafat et al., 2009; Heyes, 2016a). Continuing on this vein, I will use several well-known models and theories of individual decision-making to generate new hypotheses and derive insights on important factors influencing collective dynamics in a bottom-up manner. I will investigate how individuals incorporate social information and how this depends on the type of personal evidence individuals possess. In the following, I will introduce these models and outline their contributions to this thesis.

Models of (individual) decision-making

Multiple cue learning

Research on collectives identified diversity as a key factor for the success of groups (Krause et al., 2011; Luan et al., 2012; Herzog et al., 2019). But what makes a group diverse? A useful framework for thinking about how individuals make decisions and why individuals may arrive at different decisions is Brunswik's lens model (Brunswik, 1952). It describes situations in which a decision maker needs to judge the state of the environment, for example, whether a mushroom is toxic or not. As the state of the environment is usually not directly accessible, the individual has to infer it via observable ecological structures (i.e., cues), for example, the shape, color, or texture of the mushroom. These cues are probabilistically related to the true state of the environment (e.g., many, but not all, red mushrooms are toxic), and their validity has to be learned. Describing how individuals judge an inaccessible state of the environment through the lens of observable cues allows researchers to analyse the relationships between the environment, the cues and the judge. Extending the model to a social situation with multiple judges allows one to infer which environmental and cognitive characteristics differ between judges and, hence, cause diversity (see also Broomell and Budescu, 2009). Here, diversity refers to the degree of diverse choices (and errors) and can be influenced by (1) the diversity of cues judges observe, (2) the diversity of beliefs about the cue validities and (3) the diversity of inference strategies judges use (Fig. 1.1A). Although all factors are expected to play a role in the potential to reap the benefits of collectives, the role of belief diversity has never been explored. In Chapter 2, I investigate whether individuals develop diverse cue beliefs and whether this diversity helps individuals to profit from the collective.

Bayesian inference and heuristics

Looking at the decision-making process in a social environment through the lens of Brunswik's model also allows for a different perspective. Instead of taking the ecological structures (e.g, the color of a mushroom) as cues, one can take the choices of others as cues that need be weighted and integrated, raising the question of how individuals do so. A classic assumption is that individuals integrate multiple cues in a compensatory way, for example, via Bayesian inference or multiple regression (Fig. 1.1B). The optimal Bayesian integration of cues is achieved by weighting each (social) cue according to its reliability. Indeed, studies on perception provide evidence that individuals' decision-making process can adequately be described by Bayesian inference (Ernst and

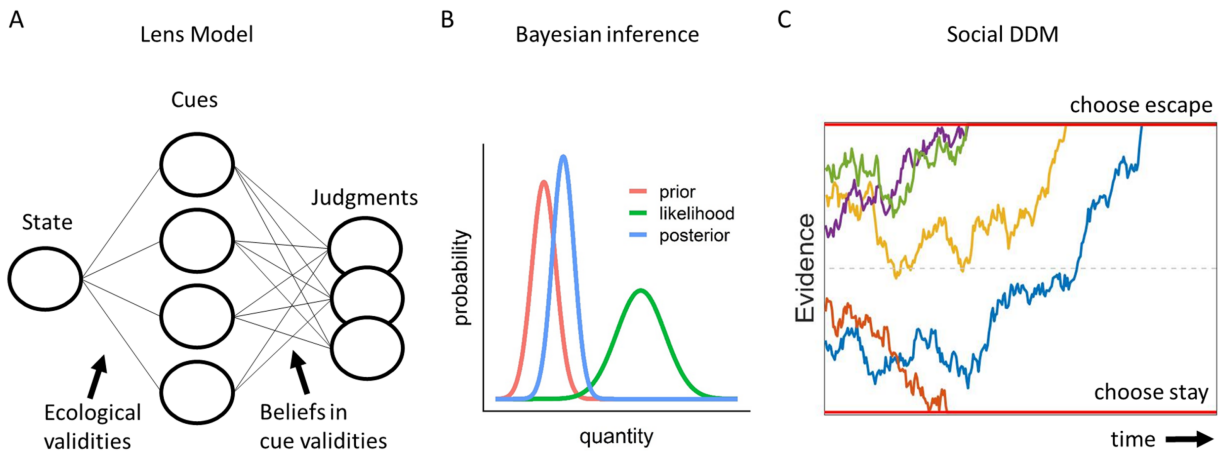


Figure 1.1. Illustrations of the three canonical individual decision-making models that will be extended to the social domain. (A) Brunswik’s Lens model (Brunswik, 1952), generalized to social environments where multiple decision makers have to judge the state of the environment. They can arrive at different conclusions (1) by observing different cues, (2) by forming different beliefs about their validity, and (3) by using diverse inference strategies. (B) Bayesian learning as an example of “optimal” cue integration. Individuals form their belief about the state of the world (e.g., a quantity) by updating their initial belief (prior) with further—possibly social—information (likelihood) to generate a new belief (posterior). (C) Illustration of the social DDM, which will be introduced in this thesis. Initially, individuals accumulate evidence independently. Individuals who begin close to either of the thresholds (red lines) are likely to choose early (by hitting the thresholds), and to provide social information to undecided individuals. This social information can influence the evidence accumulation process and sway individuals towards choice the threshold favoured by the majority.

Banks, 2002; van Dam et al., 2014). In perceptual tasks individuals usually have good information about the reliability of visual, auditory or haptic cues. Under higher levels of uncertainty, however, the feasibility of such “optimal” strategies is questioned in theories of bounded rationality. Instead, “fast and frugal” heuristics are much better suited for coping with high uncertainty (Gigerenzer and Goldstein, 1996; Gigerenzer and Gaissmaier, 2011). Accordingly, the decision maker is thought to be equipped with a *cognitive toolbox* containing a variety of heuristics or strategies to handle uncertainty. As the reliability of social cues is often highly uncertain, it is expected that individuals rely on simple heuristics when integrating social information. In Chapter 3, I will use a *cognitive toolbox* approach to investigate the repertoire of strategies individuals use to harness the wisdom of others (Rieskamp et al., 2003).

Evidence accumulation models

Another crucial element for understanding collective dynamics is time. Whether following trends (Lorenz-Spreen et al., 2019; Chen, 2008) or others in emergency evacuations (Moussaïd et al., 2016), individuals tend to imitate the choices of others, which potentially results in information

cascades (Bikhchandani et al., 1998). In such situations, individuals can often strategically time their decision. They can either decide early, thereby providing information to others, or wait until others have made their decision and use that as additional information. Individual characteristics such as knowledge (Kurvers et al., 2015; Zhang, 1997), personality (Kurvers et al., 2009) and motivation (Bousquet and Manser, 2011) can influence the timing of choice and thereby the collective dynamic. However, current models of information cascades assume random decision orders and, hence, are ill-equipped to account for decision timing (e.g., Anderson and Holt, 1997; Banerjee, 1992; Bikhchandani et al., 1998; Deneubourg et al., 1990; Mann, 2018; Sumpter and Pratt, 2008). A class of models that can address this shortcoming are the so-called drift diffusion models (DDM; Ratcliff, 1978; Ratcliff and McKoon, 2008). These models describe the decision process as a continuous updating of information over time until a decision threshold is reached and a choice is made. Consider, for example, an animal under predation risk deciding whether to escape. The animal is continuously exposed to cues potentially indicating a predator's presence (e.g., smell, sound or visual cues) and continuously updates its personal information. When enough evidence has been collected indicating the presence of a predator, the animal is expected to make the choice to escape. In Chapter 4, I will introduce a social extension of the DDM to account for choices of multiple individuals deciding sequentially (the social DDM, see Fig. 1.1C). In addition to personal information, individuals now gather social information via the choices of others who have already made a decision. Because individuals are both emitters and receivers of social information, the system is highly dynamic and accounts for the amplification of cascading social information over time. In Chapter 5 I will embed an agent-based version of the social DDM into a evolutionary algorithm to identify how individuals should optimally integrate social information and time their decision in such contexts.

Outline of the thesis

The aim of this dissertation is to deepen the knowledge of social dynamics using the above described well-known concepts of individual decision-making to cover all crucial layers of social interactions: from personal information to social learning strategies to the unfolding collective dynamics. Chapter 2 reports a study focusing on how individuals acquire and learn about the validity of cues, and how this process, in turn, affects the benefits of social information use. In the study, participants have to categorize image based on a multiple cues. The meaning of this cues has to be learned via feedback. Individuals first make a personal decision, after which they receive the decisions of all

group members and have the opportunity to revise their judgement. Simultaneously, we track the development of their beliefs in the cues. Despite observing the same images, individuals develop diverse beliefs about the validity of informative and uninformative cues. The diversity in beliefs thereby maintains error diversity in the group and, hence, promotes the quality of social information. However, participants fail to benefit from belief diversity because they only incorporate opposing views if this view is supported by a very large majority (i.e., if the group members show strong agreement). Although the participants increase their sensitivity to social information over the course of the experiment, they maintain their reluctance to follow small majorities. Using simulations I show that this reliance on strongly agreeing majorities impedes individuals from benefiting from diversity. I conclude that diversity is an important factor promoting the quality of social information but also highlight how social learning strategies can prevent individuals from exploiting these benefits.

Chapter 3 provides a more detailed picture of how people cope with disagreement. In the study, individuals are confronted with an estimation task and allowed to adjust their initial judgment after observing judgements of other peers. Controlling the exact distribution and distance of the social information allowed me to decouple characteristics of personal and social information. I investigate individuals' behavior with a cognitive model encompassing a toolbox of strategies, including simple heuristics and more sophisticated Bayesian updating. I conclude that the incorporation of social information strongly depends on the similarity to one's own judgment and degree of agreement among others. Additionally, individuals strongly differ in the heuristics they use incorporate social information. The results help elucidate what factors are conducive to people changing their minds. This contributes to the understanding of how individuals' social information use shapes the spread of information and opinions in our interconnected society.

Providing social information as a collection of independent judgments is a rather simplistic process of social information exchange. Decisions are usually not made simultaneously but take place over time allowing information to spread via *information cascades*. These more realistic situations confront individuals with new challenges. Should one decide quickly or wait for further information, and how should social information which is changing over time be weighted? In Chapter 4, I introduce the social DDM to investigate the process of how individuals in groups integrate personal and social information dynamically over time. I test the model in a sequential choice paradigm in which participants need to decide between two options and are free to time their decision. I find that correct information spreads when individuals' confidence accurately

reflects their personal information and when more confident individuals decide faster. Under these conditions, late-deciding (unconfident) individuals are more likely to adopt the social information, amplifying the correct signal provided by early-deciding individuals. The social DDM successfully captures all the key dynamics of the social system, thereby revealing the cognitive mechanisms underpinning information cascades.

Many choice environments are characterised by asymmetric errors (e.g., missing the presence of a predator or a malignant tumour is worse than a false alarm). Under conditions in which errors are asymmetrical, individuals should develop response biases to avoid the more costly error. However, the consequences of such asymmetric error costs on group dynamics are poorly understood. In Chapter 5, I embed the social DDM into an evolutionary algorithm to derive the optimal behaviour in sequential decision tasks under asymmetric error costs. The evolutionary simulations show that single individuals should, indeed, develop response biases to avoid the worse outcome. However, the presence of such response biases in groups can be extremely detrimental, because in large groups biases can quickly amplify, triggering false information cascades. As a result, individuals in large groups should use much weaker response biases to benefit from social information. I show that individuals facing asymmetric error costs in social environments need to carefully trade off the expressed bias and sensitivity to social information to avoid large error costs, while simultaneously benefit from the collective.

Lastly, Chapter 6 provides a synthesis of the findings embedding the main results in the broader context of social learning and collective dynamics, and proposes directions of future research.

References

- Abrahamson, E. and Rosenkopf, L. (1997). Social network effects on the extent of innovation diffusion: A computer simulation. *Organization Science*, 8(3):289–309.
- Anderson, L. R. and Holt, C. A. (1997). Information cascades in the laboratory. *The American Economic Review*, 87(5):847–862.
- Aplin, L. M., Farine, D. R., Morand-Ferron, J., Cockburn, A., Thornton, A., and Sheldon, B. C. (2015). Experimentally induced innovations lead to persistent culture via conformity in wild birds. *Nature*, 518(7540):538–541.
- Asch, S. E. (1956). Studies of independence and conformity: I. a minority of one against a unanimous majority. *Psychological Monographs: General and Applied*, 70(9):1–70.
- Asch, S. E. and Guetzkow, H. (1951). *Effects of group pressure upon the modification and distortion of judgments*. Carnegie Press, Oxford, England.
- Banerjee, A. V. (1992). A simple model of herd behavior. *The Quarterly Journal of Economics*, 107(3):797–817.
- Bang, D. and Frith, C. D. (2017). Making better decisions in groups. *Royal Society Open Science*, 4(8):170193.
- Battaglini, M. (2005). Sequential voting with abstention. *Games and Economic Behavior*, 51(2):445–463.

- Baude, M., Dajoz, I., and Danchin, E. (2008). Inadvertent social information in foraging bumblebees: effects of flower distribution and implications for pollination. *Animal Behaviour*, 76(6):1863–1873.
- Bikhchandani, S., Hirshleifer, D., and Welch, I. (1998). Learning from the Behavior of Others: Conformity, Fads, and Informational Cascades. *Journal of Economic Perspectives*, 12(3):151–170.
- Biro, D., Sumpter, D. J., Meade, J., and Guilford, T. (2006). From compromise to leadership in pigeon homing. *Current Biology*, 16(21):2123–2128.
- Bond, R. and Smith, P. B. (1996). Culture and conformity: A meta-analysis of studies using asch’s (1952b, 1956) line judgment task. *Psychological Bulletin*, 119(1):111–137.
- Bond, R. M., Fariss, C. J., Jones, J. J., Kramer, A. D., Marlow, C., Settle, J. E., and Fowler, J. H. (2012). A 61-million-person experiment in social influence and political mobilization. *Nature*, 489(7415):295–298.
- Bousquet, C. A. and Manser, M. B. (2011). Resolution of experimentally induced symmetrical conflicts of interest in meerkats. *Animal Behaviour*, 81(6):1101–1107.
- Boyd, R. and Richerson, P. J. (1985). *Culture and the evolutionary process*. University of Chicago Press.
- Boyd, R. and Richerson, P. J. (1995). Why does culture increase human adaptability? *Ethology and Sociobiology*, 16(2):125–143.
- Boyd, R. and Richerson, P. J. (1996). Why culture is common, but cultural evolution is rare. *Proceedings of the British Academy*, 88:77–94.
- Bridges, A. D. and Chittka, L. (2019). Animal behaviour: Conformity and the beginnings of culture in an insect. *Current Biology*, 29(5):R167 – R169.
- Broomell, S. B. and Budescu, D. V. (2009). Why are experts correlated? decomposing correlations between judges. *Psychometrika*, 74(3):531–553.
- Brunswik, E. (1952). The conceptual framework of psychology. *Psychological Bulletin*, 49(6):654–656.
- Chamley, C. and Gale, D. (1994). Information revelation and strategic delay in a model of investment. *Econometrica*, 62(5):1065.
- Chen, Y.-F. (2008). Herd behavior in purchasing books online. *Computers in Human Behavior*, 24(5):1977–1992.
- Condorcet, M. J. A. N. d. C. (1785). *Essai sur l’application de l’analyse à la probabilité des décisions rendues à la pluralité des voix*. Imprimerie Royale, Paris, France.
- Couzin, I. D. and Krause, J. (2003). Self-organization and collective behavior in vertebrates. In Slater, P., Rosenblatt, J., and Roper, T., editors, *Advances in the Study of Behavior*. Academic Press, London, England.
- Couzin, I. D., Krause, J., Franks, N. R., and Levin, S. A. (2005). Effective leadership and decision-making in animal groups on the move. *Nature*, 433(7025):513.
- Danchin, E., Giraldeau, L.-A., Valone, T. J., and Wagner, R. H. (2004). Public information: from nosy neighbors to cultural evolution. *Science*, 305(5683):487–491.
- Day, R. L., MacDonald, T., Brown, C., Laland, K. N., and Reader, S. M. (2001). Interactions between shoal size and conformity in guppy social foraging. *Animal Behaviour*, 62(5):917–925.
- Deaner, R. O., Khera, A. V., and Platt, M. L. (2005). Monkeys pay per view: adaptive valuation of social images by rhesus macaques. *Current Biology*, 15(6):543–548.
- Deneubourg, J.-L. (1977). Application de l’ordre par fluctuations a la description de certaines étapes de la construction du nid chez les termites. *Insectes Sociaux*, 24(2):117–130.
- Deneubourg, J.-L., Aron, S., Goss, S., and Pasteels, J. M. (1990). The self-organizing exploratory pattern of the Argentine ant. *Journal of Insect Behavior*, 3(2):159–168.
- Deutsch, M. and Gerard, H. B. (1955). A study of normative and informational social influences upon individual judgment. *The Journal of Abnormal and Social Psychology*, 51(3):629–636.
- Dindo, M., Whiten, A., and de Waal, F. B. (2009). In-group conformity sustains different foraging traditions in capuchin monkeys (*cebus apella*). *PLoS ONE*, 4(11):e7858.
- Drea, C. M. and Wallen, K. (1999). Low-status monkeys “play dumb” when learning in mixed social groups. *Proceedings of the National Academy of Sciences*, 96(22):12965–12969.
- Ernst, M. O. and Banks, M. S. (2002). Humans integrate visual and haptic information in a statistically optimal fashion. *Nature*, 415(6870):429–433.
- Feldman, M. W., Aoki, K., and Kumm, J. (1996). Individual versus social learning: evolutionary analysis in a fluctuating environment. *Anthropological Science*, 104(3):209–231.

- Flack, A., Biro, D., Guilford, T., and Freeman, R. (2015). Modelling group navigation: transitive social structures improve navigational performance. *Journal of The Royal Society Interface*, 12(108):20150213.
- Fujisaki, I., Honda, H., and Ueda, K. (2018). Diversity of inference strategies can enhance the ‘wisdom-of-crowds’ effect. *Palgrave Communications*, 4(1):107.
- Galef, B. G. (2009). Strategies for social learning: testing predictions from formal theory. *Advances in the Study of Behavior*, 39:117–151.
- Galef, B. G. and Whiskin, E. E. (2008). Conformity in norway rats? *Animal Behaviour*, 75(6):2035–2039.
- Germar, M., Schlemmer, A., Krug, K., Voss, A., and Mojzisch, A. (2014). Social influence and perceptual decision making: A diffusion model analysis. *Personality and Social Psychology Bulletin*, 40(2):217–231.
- Gigerenzer, G. and Gaissmaier, W. (2011). Heuristic decision making. *Annual Review of Psychology*, 62:451–482.
- Gigerenzer, G. and Goldstein, D. G. (1996). Reasoning the fast and frugal way: Models of bounded rationality. *Psychological Review*, 103(4):650–669.
- Giraldeau, L., Valone, T. J., and Templeton, J. J. (2002). Potential disadvantages of using socially acquired information. *Philosophical Transactions of the Royal Society of London. Series B: Biological Sciences*, 357(1427):1559–1566.
- Goldstone, R. L. and Gureckis, T. M. (2009). Collective behavior. *Topics in Cognitive Science*, 1(3):412–438.
- Granovetter, M. (1978). Threshold models of collective behavior. *American Journal of Sociology*, 83(6):1420–1443.
- Gul, F. and Lundholm, R. (1995). Endogenous timing and the clustering of agents’ decisions. *Journal of Political Economy*, 103(5):1039–1066.
- Hastie, R. and Kameda, T. (2005). The robust beauty of majority rules in group decisions. *Psychological review*, 112(2):494–508.
- Helbing, D., Farkas, I. J., Molnar, P., and Vicsek, T. (2002). Simulation of pedestrian crowds in normal and evacuation situations. *Pedestrian and Evacuation Dynamics*, 21(2):21–58.
- Henrich, J. and Boyd, R. (1998). The evolution of conformist transmission and the emergence of between-group differences. *Evolution and Human Behavior*, 19(4):215–241.
- Herzog, S. M., Litvinova, A., Yahosseini, K. S., Tump, A. N., and Kurvers, R. H. (2019). The ecological rationality of the wisdom of crowds. In *Taming Uncertainty*. MIT Press.
- Heyes, C. (2016a). Blackboxing: social learning strategies and cultural evolution. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 371(1693):20150369.
- Heyes, C. (2016b). Who knows? metacognitive social learning strategies. *Trends in Cognitive Sciences*, 20(3):204–213.
- Hoogendoorn, S. P. and Daamen, W. (2007). Microscopic calibration and validation of pedestrian models: Cross-comparison of models using experimental data. In Schadschneider, A., Pöschel, T., Kühne, R., Schreckenberg, M., and Wolf, D. E., editors, *Traffic and Granular Flow’05*, pages 329–340. Springer Berlin Heidelberg, Berlin, Heidelberg.
- Horner, V., Proctor, D., Bonnie, K. E., Whiten, A., and de Waal, F. B. (2010). Prestige affects cultural learning in chimpanzees. *PloS ONE*, 5(5):e10625.
- Kaniovski, S. (2010). Aggregation of correlated votes and condorcet’s jury theorem. *Theory and Decision*, 69(3):453–468.
- Kawaguchi, L. G., Ohashi, K., and Toquenaga, Y. (2007). Contrasting responses of bumble bees to feeding conspecifics on their familiar and unfamiliar flowers. *Proceedings of the Royal Society B: Biological Sciences*, 274(1626):2661–2667.
- Kempe, M. and Mesoudi, A. (2014). Experimental and theoretical models of human cultural evolution. *Wiley Interdisciplinary Reviews: Cognitive Science*, 5(3):317–326.
- Kendal, R. L., Coolen, I., and Laland, K. N. (2004). The role of conformity in foraging when personal and social information conflict. *Behavioral Ecology*, 15(2):269–277.
- Konopasky, R. J. and Telegdy, G. A. (1977). Conformity in the rat: A leader’s selection of door color versus a learned door-color discrimination. *Perceptual and Motor Skills*, 44(1):31–37.
- Krause, J., Ruxton, G. D., and Krause, S. (2010). Swarm intelligence in animals and humans. *Trends in Ecology & Evolution*, 25(1):28–34.
- Krause, S., James, R., Faria, J. J., Ruxton, G. D., and Krause, J. (2011). Swarm intelligence in humans:

- diversity can trump ability. *Animal Behaviour*, 81(5):941 – 948.
- Kurvers, R. H., Eijkelenkamp, B., van Oers, K., van Lith, B., van Wieren, S. E., Ydenberg, R. C., and Prins, H. H. (2009). Personality differences explain leadership in barnacle geese. *Animal Behaviour*, 78(2):447–453.
- Kurvers, R. H., Wolf, M., Naguib, M., and Krause, J. (2015). Self-organized flexible leadership promotes collective intelligence in human groups. *Royal Society Open Science*, 2(12):150222.
- Laan, A., Madirolas, G., and de Polavieja, G. G. (2017). Rescuing Collective Wisdom when the Average Group Opinion Is Wrong. *Frontiers in Robotics and AI*, 4:56.
- Ladha, K. K. (1992). The Condorcet Jury Theorem, free speech, and correlated votes. *American Journal of Political Science*, 36(3):617–634.
- Ladha, K. K. (1995). Information pooling through majority-rule voting: Condorcet’s jury theorem with correlated votes. *Journal of Economic Behavior and Organization*, 26(3):353–372.
- Laland, K. N. (2004). Social learning strategies. *Animal Learning & Behavior*, 32(1):4–14.
- Latané, B. (1981). The psychology of social impact. *American psychologist*, 36(4):343–356.
- Lewandowsky, S., Ecker, U. K., and Cook, J. (2017). Beyond misinformation: Understanding and coping with the “post-truth” era. *Journal of Applied Research in Memory and Cognition*, 6(4):353–369.
- List, C., Elsholtz, C., and Seeley, T. D. (2009). Independence and interdependence in collective decision making: an agent-based model of nest-site choice by honeybee swarms. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 364(1518):755–762.
- Lorenz, J. (2007). Continuous opinion dynamics under bounded confidence: A survey. *International Journal of Modern Physics C*, 18(12):1819–1838.
- Lorenz-Spreen, P., Mønsted, B. M., Hövel, P., and Lehmann, S. (2019). Accelerating dynamics of collective attention. *Nature Communications*, 10(1):1759.
- Luan, S., Katsikopoulos, K. V., and Reimer, T. (2012). When does diversity trump ability (and vice versa) in group decision making? A simulation study. *PLoS ONE*, 7(2):e31043.
- Mann, R. P. (2018). Collective decision making by rational individuals. *Proceedings of the National Academy of Sciences*, 115(44):E10387–E10396.
- Mannes, A. E., Soll, J. B., and Larrick, R. P. (2014). The wisdom of select crowds. *Journal of personality and social psychology*, 107(2):276.
- McElreath, R., Lubell, M., Richerson, P. J., Waring, T. M., Baum, W., Edsten, E., Efferson, C., and Paciotti, B. (2005). Applying evolutionary models to the laboratory study of social learning. *Evolution and Human Behavior*, 26(6):483–508.
- Milgram, S., Bickman, L., and Berkowitz, L. (1969). Note on the drawing power of crowds of different size. *Journal of Personality and Social Psychology*, 13(2):79–82.
- Morgan, T. and Laland, K. (2012). The biological bases of conformity. *Frontiers in Neuroscience*, 6:87.
- Morgan, T. J. H., Rendell, L. E., Ehn, M., Hoppitt, W., and Laland, K. N. (2012). The evolutionary basis of human social learning. *Proceedings of the Royal Society B: Biological Sciences*, 279(1729):653–662.
- Moussaïd, M., Kämmer, J. E., Analytis, P. P., and Neth, H. (2013). Social Influence and the Collective Dynamics of Opinion Formation. *PLoS ONE*, 8(11):e78433.
- Moussaïd, M., Kapadia, M., Thrash, T., Sumner, R. W., Gross, M., Helbing, D., and Hölscher, C. (2016). Crowd behaviour during high-stress evacuations in an immersive virtual environment. *Journal of The Royal Society Interface*, 13(122):20160414.
- Pike, T. W. and Laland, K. N. (2010). Conformist learning in nine-spined sticklebacks’ foraging decisions. *Biology letters*, 6(4):466–468.
- Raafat, R. M., Chater, N., and Frith, C. (2009). Herding in humans. *Trends in Cognitive Sciences*, 13(10):420–428.
- Ratcliff, R. (1978). A theory of memory retrieval. *Psychological Review*, 85(2):59–108.
- Ratcliff, R. and McKoon, G. (2008). The diffusion decision model: theory and data for two-choice decision tasks. *Neural Computation*, 20(4):873–922.
- Reebs, S. G. (2001). Influence of body size on leadership in shoals of golden shiners, *notemigonus crysoleucas*. *Behaviour*, 138(7):797–809.
- Rieskamp, J., Busemeyer, J. R., and Laine, T. (2003). How Do People Learn to Allocate Resources? Comparing Two Learning Theories. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 29(6):1066–1081.

- Rogers, A. R. (1988). Does Biology Constrain Culture? *American Anthropologist*, 90(4):819–831.
- Rogers, E. M. (2004). A prospective and retrospective look at the diffusion model. *Journal of Health Communication*, 9(sup1):13–19. PMID: 14960401.
- Salganik, M. J., Dodds, P. S., and Watts, D. J. (2006). Experimental study of inequality and unpredictability in an artificial cultural market. *Science*, 311(5762):854–856.
- Schlag, K. H. (1998). Why Imitate, and If So, How?: A Boundedly Rational Approach to Multi-armed Bandits. *Journal of Economic Theory*, 78(1):130–156.
- Sekiguchi, T. and Ohtsuki, H. (2015). Effective group size of majority vote accuracy in sequential decision-making. *Japan Journal of Industrial and Applied Mathematics*, 32(3):595–614.
- Shiller, R. J. (2002). Bubbles, human judgment, and expert opinion. *Financial Analysts Journal*, 58(3):18–26.
- Slutkin, G., Ransford, C., and Decker, R. B. (2015). Cure violence: Treating violence as a contagious disease. In Maltz, M. D., editor, *Envisioning Criminology: Researchers on Research as a Process of Discovery*, pages 43–56. Springer International Publishing, Cham.
- Smolla, M., Alem, S., Chittka, L., and Shultz, S. (2016). Copy-when-uncertain: bumblebees rely on social information when rewards are highly variable. *Biology Letters*, 12(6):20160188.
- Sorkin, R., Hays, C., and West, R. (2001). Signal-detection analysis of group decision making. *Psychological Review*, 108(1):183–203.
- Strandburg-Peshkin, A., Twomey, C. R., Bode, N. W., Kao, A. B., Katz, Y., Ioannou, C. C., Rosenthal, S. B., Torney, C. J., Wu, H. S., Levin, S. A., and Couzin, I. D. (2013). Visual sensory networks and effective information transfer in animal groups. *Current Biology*, 23(17):R709 – R711.
- Sumpter, D. J. and Pratt, S. C. (2008). Quorum responses and consensus decision making. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 364(1518):743–753.
- Surowiecki, J. (2004). *The Wisdom of Crowds: Why the Many Are Smarter Than the Few and How Collective Wisdom Shapes Business, Economies, Societies and Nations*. Doubleday, New York.
- Toelch, U., Bruce, M. J., Newson, L., Richerson, P. J., and Reader, S. M. (2014). Individual consistency and flexibility in human social information use. *Proceedings of the Royal Society B: Biological Sciences*, 281(1776):20132864.
- Toelch, U. and Dolan, R. J. (2015). Informational and normative influences in conformity from a neurocomputational perspective. *Trends in cognitive sciences*, 19(10):579–589.
- Toelch, U., Panizza, F., and Heekeren, H. R. (2018). Norm compliance affects perceptual decisions through modulation of a starting point bias. *Royal Society Open Science*, 5(3):171268.
- Toyokawa, W., Whalen, A., and Laland, K. N. (2019). Social learning strategies regulate the wisdom and madness of interactive crowds. *Nature Human Behaviour*, 3(2):183.
- van Dam, L. C., Parise, C. V., and Ernst, M. O. (2014). Modeling multisensory integration. In *Sensory integration and the unity of consciousness*, page 209. MIT press.
- Van Vugt, M. (2006). Evolutionary origins of leadership and followership. *Personality and Social Psychology Review*, 10(4):354–371.
- Vosoughi, S., Roy, D., and Aral, S. (2018). The spread of true and false news online. *Science*, 359(6380):1146–1151.
- Webster, M. and Laland, K. (2008). Social learning strategies and predation risk: minnows copy only when using private information would be costly. *Proceedings of the Royal Society B: Biological Sciences*, 275(1653):2869–2876.
- Welch, I. (2000). Herding among security analysts. *Journal of Financial Economics*, 58(3):369–396.
- Whiten, A., Horner, V., and De Waal, F. B. (2005). Conformity to cultural norms of tool use in chimpanzees. *Nature*, 437(7059):737–740.
- Xiong, F. and Liu, Y. (2014). Opinion formation on social media: An empirical approach. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 24(1):013130.
- Yaniv, I. and Kleinberger, E. (2000). Advice taking in decision making: Egocentric discounting and reputation formation. *Organizational Behavior and Human Decision Processes*, 83(2):260 – 281.
- Zhang, J. (1997). Strategic delay and the onset of investment cascades. *The RAND Journal of Economics*, 28(1):188–205.
- Ziegelmeier, A., My, K. B., Vergnaud, J.-C., Willinger, M., et al. (2005). Strategic delay and rational imitation in the laboratory. Unpublished, Max-Planck-Inst. for Research into Economic Systems.

Chapter 2

Individuals fail to reap the collective benefits of diversity because of over-reliance on personal information

Tump, A.N., Wolf, M., Krause, J., & Kurvers, R.H.J.M. (2018).
Journal of the Royal Society Interface, 15(142), 20180155.

This chapter can be accessed at:

<https://doi.org/10.1098/rsif.2018.0155>

Chapter 3

Strategies for integrating disparate social information

Molleman, L.* , Tump, A.N.* , Gradassi, A., Jayles, B., Herzog, S.,
Kurvers, R.H.J.M. & van den Bos, W.

In Preparation

*These authors contributed equally

Abstract

Our increasingly interconnected world provides virtually unlimited opportunities to observe the behavior of others. This affords abundant useful information but also requires navigating complex social environments with people holding disparate or conflicting views. It is, however, still largely unresolved which strategies people use to integrate information from multiple social sources that (dis)agree with oneself and each other and how this affects collective dynamics. We address this in three steps. First, systematically varying the variance and skewness of the social information in a highly controlled experiment on social influence, we show people’s use of social information strongly depends on how it is distributed. We found that, as expected, higher variance among social sources reduces their social influence. More importantly, only observing one social source confirming an individual’s estimate, resulted in a strong decrease of influence of other—more distant—social sources. Second, we develop a framework for modelling the cognitive processes underlying the integration of disparate social information, combining Bayesian inference with heuristic approaches. Our models accurately account for people’s adjustment strategies and reveal that people particularly heed social information that confirms personal judgments. Moreover, we find strong inter-individual differences in strategy use. Third, using simulations across a range of opinion distributions in virtual groups, we provide novel insights into how identified adjustment strategies can promote the emergence of filter bubbles and exacerbate group polarization. Overall, our results help elucidate what aspects of the social environment are, and are not, conducive to changing people’s minds.

Acknowledgments: We thank Casper Hesp and the members of the Connected Minds Lab at the University of Amsterdam for useful comments and discussions. L.M. and W.v.d.B. were supported by Open Research Area (ID 176); L.M. is further supported by an Amsterdam Brain and Cognition Project grant 2018. W.v.d.B. is further supported by the Jacobs Foundation, European Research Council grant (ERC-2018-StG-803338) and the Netherlands Organization for Scientific Research grant (NWO-VIDI 016.Vidi.185.068). R.H.J.M.K. acknowledges funding from the German Research Foundation (grant number: KU 3369/1-1).

Author Contributions: L.M., A.G., W.v.d.B. designed the study and collected data; A.N.T., and L.M. analyzed data; A.N.T., and L.M., carried out the simulations; L.M., A.N.T., R.H.J.M.K. and W.v.d.B. drafted the manuscript with substantial input from A.G., B.J. and S.H..

Introduction

The way in which individuals integrate personal and social information shapes a wide range of collective phenomena such as the spread of knowledge across social networks (Bakshy et al., 2012), the development of financial markets (Devenow and Welch, 1996), political mobilization (Bloomfield and Hales, 2009) and voting outcomes (Stewart et al., 2019). By interacting with others and observing their behavior, individuals can often glean useful information helping them to, for example, rapidly adjust to new environments (Boyd et al., 2011). However, with more and more of these interactions occurring online, and being controlled by algorithms that prioritize interactions between like-minded people, people can become cut off from information that might challenge their beliefs (Sunstein, 2007; Pariser, 2011). It has been argued that such dynamics may lead to stronger ingroup consensus, and between-group polarization, in controversial matters ranging from same-sex marriage and vaccine safety to climate change (Lewandowsky et al., 2013; Deryugina and Shurchkov, 2016; Kerr and Wilson, 2018). To counteract these dynamics, several recent projects (e.g., onesub.io and nuzzera.com) have focused on de-biasing news feeds and providing users with more balanced social information from disparate sources (for review see Bozdag and van den Hoven, 2015). Yet, to effectively de-bias individuals we need a detailed understanding of the strategies individuals use when they are confronted with social information coming from multiple sources.

Social information use has been extensively studied across the biological and social sciences (Asch, 1956; Boyd and Richerson, 1985; Surowiecki, 2004; Bond, 2005; Page, 2008; Rendell et al., 2010; Mesoudi, 2011; Hoppitt and Laland, 2013; Kurvers et al., 2014; van den Berg et al., 2015; Aplin et al., 2015; Moussaïd et al., 2017; Danchin et al., 2018; Kendal et al., 2018; Analytis et al., 2018; Tump et al., 2018; Derex et al., 2019; Mercier and Morin, 2019; Efferson et al., 2019). In humans, social information use often involves changing one's mind after observing the behavior of other people (Yaniv, 1997, 2004; Soll and Larrick, 2009; Bednarik and Schultze, 2015). This process is commonly emulated with estimation tasks in which people are allowed to revise their first estimates after observing the estimate of a peer (e.g., Yaniv and Kleinberger, 2000; Bonaccio and Dalal, 2006; Moussaïd et al., 2013; Moussaïd et al., 2017; Jayles et al., 2017). Studies using this approach give a detailed and quantified account of the effects of social cues on behavior, primarily focusing on how individuals incorporate a single piece of social information (Yaniv, 1997, 2004; Bonaccio and Dalal, 2006; Soll and Larrick, 2009; Bednarik and Schultze, 2015; Molleman et al.,

2019). Studies considering multiple peers have mainly evaluated the effect of the central tendency (the mean social information; Larrick and Soll, 2006; Park et al., 2017; Jayles et al., 2017; but see Yaniv and Milyavsky, 2007). In most real-world environments, however, people are confronted with multiple sources of social information simultaneously with varying degrees of extremity. Currently it is unclear how people integrate such disparate social information. Here we will address this important issue in three steps.

First, we will experimentally investigate how key characteristics of the distribution of social information shape social information use. Specifically, we systematically vary the variance (reflecting the agreement among peers) and skewness (reflecting the clustering of peers close or far away from the focal participant) of the distribution, while holding its mean constant. We show that the impact of social information strongly depends on how it is distributed. Disagreement among peers decreased their overall influence. Furthermore, the direction of the skew substantially altered the impact of social information: participants adjusted their first estimate considerably more when the majority of peers moderately agreed with them and one peer strongly disagreed, compared to a situation in which a single peer strongly confirmed them, but the majority of peers strongly disagreed. This highlights strong effects of confirmation bias.

Second, we will introduce a formal model to explain the strategies underlying these adjustments. When observing information from a single social source, individuals have been shown to use different strategies: (1) keeping one's initial beliefs, (2) adopting the behavior of others, or (3) 'compromising' between personal and social information (Yaniv, 1997; Budescu et al., 2003; Budescu and Yu, 2007; Harries et al., 2004; Bahrami et al., 2010; Shea et al., 2014; Aitchison et al., 2015; Toelch and Dolan, 2015; Soll and Larrick, 2009; Yaniv, 2004; Moussaïd et al., 2013; Moussaïd et al., 2017). When compromising, people generally take a (weighted) average of their initial beliefs and the social information, weighting their own beliefs more than the social information, and weighting social information more if it confirms their own beliefs (Yaniv and Kleinberger, 2000; Yaniv, 2004). Which adjustment strategies people use when facing multiple pieces of disparate social information is, however, largely unknown. To address this, we develop novel cognitive models that provide a unifying theoretical framework that extends previous modeling efforts (Toelch and Dolan, 2015; Bang and Frith, 2017; Park et al., 2017), accommodating both simple heuristics (e.g., keeping and adopting) and more complex strategies (e.g., compromising). Our modelling results reveal that social information receives more weight when it is confirming the participant's initial belief and when it is in close agreement with other social information (reflecting peer consensus).

Moreover, our models accurately predict how the distribution of social information impacts the relative frequencies of adjustment strategies, providing a unified account for when people are likely to keep their own beliefs, adopt the behavior of a peer, or compromise towards the mean social information.

Finally, we use our model to generate predictions for social information use beyond the scenarios studied in our experiment. These simulations reveal that even in making ‘informational’ decisions with minimal social context, people’s prioritizing of social information that confirm their personal beliefs can exacerbate filter bubble effects and can even lead some individuals to adopt more extreme beliefs, fueling group polarization.

Experimental Design

To examine how people integrate disparate information from multiple social sources, we used an adapted version of the BEAST (Berlin Estimate AdjuStment Task): a validated perceptual judgment task known to reliably measure individuals’ social information use (Fig. 3.1; Molleman et al., 2019). In the task, participants are repeatedly shown images of animal groups and have to estimate the number of animals (Fig. 3.1A,B). They then observe the estimates of three previous participants, and make a second estimate (Fig. 3.1C). The relative degree of adjustment (s) is a measure of an individual’s social information use (Fig 3.1D; Molleman et al., 2019).

We studied participants’ social information use across four conditions that systematically differed in variance and skewness, while controlling for the mean deviation from a participant’s first estimate (Fig. 3.1E): (*i*) low variance, not skewed (LN); (*ii*) high variance, not skewed (HN); (*iii*) high variance, with the distribution leaning away from the participants’ first estimate (HA; one peer strongly agrees with the participant, and two peers strongly disagree); (*iv*) high variance, with the distribution leaning towards the participant’s first estimate (HT; two peers show moderate agreement, and one shows strong disagreement). These conditions encompass a broad range of distributions an individual may encounter when sampling its social environment. Across all conditions, the three pieces of social information always pointed in the same-and correct-direction (i.e., avoiding situations in which the social information bracketed the personal estimate). Importantly, holding constant the mean deviation from a participant’s first estimate across conditions implies that a participant weighting all peer estimates equally should make similar adjustments across all conditions.

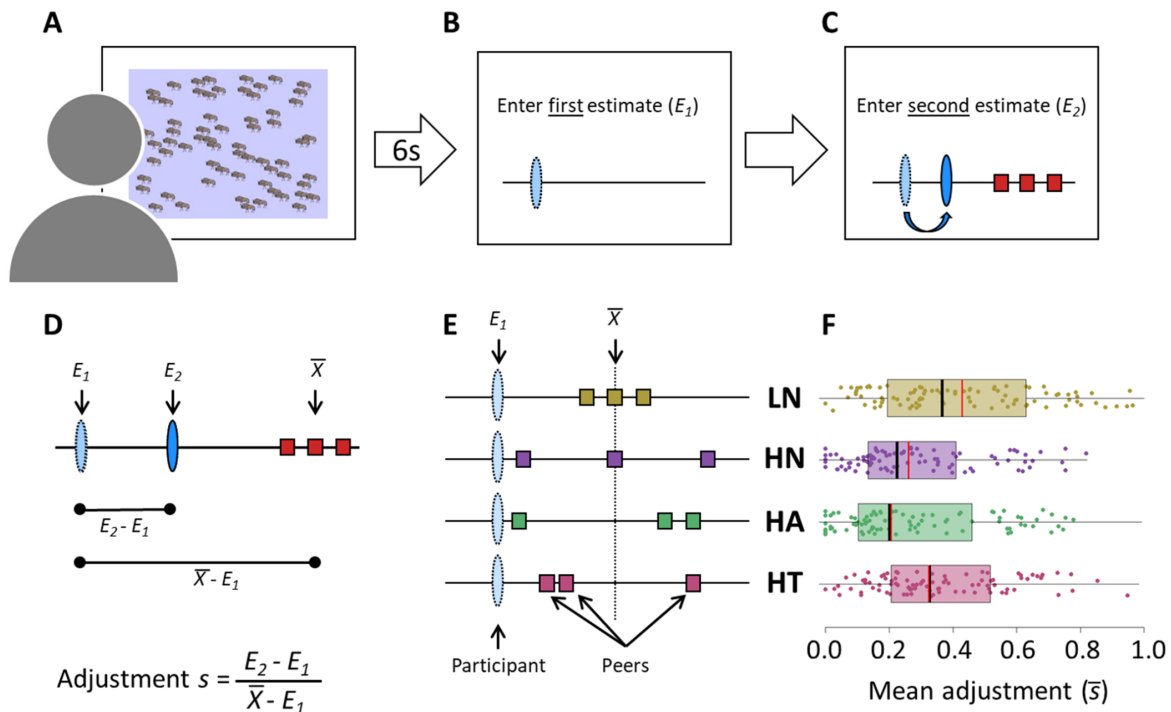


Figure 3.1. Experimental paradigm and the impact of disparate social information. (A) Participants start with observing a group of animals for six seconds. (B) Next, they enter their first estimate of the total number of animals using a slider. (C) Then, they observe the estimates of three pre-recorded peers (red squares), as well as their first estimate (light blue oval), and enter their second estimate (dark blue oval). (D) Social information use in a round (s) is calculated as the adjustment from the first estimate (E_1) to the second estimate (E_2), divided by the distance between the first estimate and the mean of the social information (\bar{X}). (E) In four conditions, we manipulated the distribution of social information (squares) relative to a participant’s first estimate (oval). Across conditions, we varied the variance and skewness of the social information, while fixing the distance between the mean of the social information (\bar{X}) and the participant’s first estimate (E_1 ; for details see Experimental Design). Peer estimates displayed either low variance (LN) or high variance - but no skewness (HN) -, or high variance with the distribution leaning away from (HA), or towards (HT) E_1 . (F) Mean estimate shifts in each condition. Colored dots show participants’ mean adjustments across the five rounds of each treatment; $\bar{s} = 1$ indicates a mean estimate shifts to \bar{X} . Boxplots show the interquartile range (IQR), the median (black line), and the 1.5 IQR (whiskers). Red vertical lines show predicted medians of the best-fitting model for each treatment (see Cognitive Model).

Ninety-five participants completed thirty rounds of the judgment task. Each participant completed five rounds of each condition, presented in randomized order and interleaved with 10 ‘filler’ rounds. Social information in the ‘filler’ rounds consisted of the estimates of three randomly drawn previous participants, ensuring that across all rounds, social information was sometimes higher and sometimes lower than a participant’s first estimate, and sometimes bracketed the participant’s personal estimate (see Methods for details). As a control, participants completed five rounds in which they did not observe the stimulus themselves, but only the estimates of four peers. The distribution of these peer estimates emulated the distributions of the four experimental conditions (i.e.,

one of each condition), plus the filler round. This enables us to compare how participants integrate personal and social information with a control in which participants integrate four pieces of social information only. Participants were recruited and performed the task online (see Methods).

Results

Experimental results. Participants' use of social information strongly depended on its distribution (Fig. 3.1F). Participants adjusted their estimates most when social information had low variance and no skewness (LN condition; Fig. 3.1F, yellow), shifting, on average, 41.5% towards the mean social information. In the high variance and no skew condition (HN condition; Fig. 3.1F, purple), average adjustments were credibly lower (mean adjustment: 29.0%; see Table B1 for statistics). Although adjustments in both conditions with skewed distributions were credibly lower than in the LN condition, the direction of the skew affected the relative adjustment: participants adjusted credibly more when the distribution of social information leaned towards the participants (HT condition; mean: 36.5%; Fig. 3.1F, red) than when it leaned away from the participant (HA condition; mean: 27.9%; Fig. 3.1F, green). Interestingly, we observed strong correlations between participants' mean adjustments across conditions (all pairwise Pearson correlations $r \leq 0.76, P < 0.001$; Fig. B1), indicating strong inter-individual differences in social information use (see also below).

Figure 3.2 zooms in on the strategies underlying behavioral adjustments across rounds, differentiating between three distinct strategies: (1) *keeping* the first estimate, (2) *adopting* the estimate of one of the three peers, or (3) *compromising* between the first estimate and the peer estimates. The relative frequency of these strategies differed between the four conditions (Fig. 3.2A-D; see Table B2 for statistics). When participants observed a single peer that closely agreed with them (i.e., the HN and HA conditions; Fig. 3.2B,C), participants were more likely to (i) either keep their first estimate (ii) or adopt the estimate of this close peer, and (iii) less likely to compromise. As a consequence, participants adjusted less in these conditions. Figure 3.2E shows the frequency of strategies per participant across all conditions, illustrating that participants ranged from almost exclusively compromising, to exclusively keeping (see also Fig. B1), with compromising being the most frequent strategy.

In all four control conditions, in which participants did not observe the stimulus—but four peer estimates, emulating the four distributions of the experimental conditions—, responses were close to the arithmetic mean of the four peer estimates (Fig. B2; Fig. B3). Participants did, however,

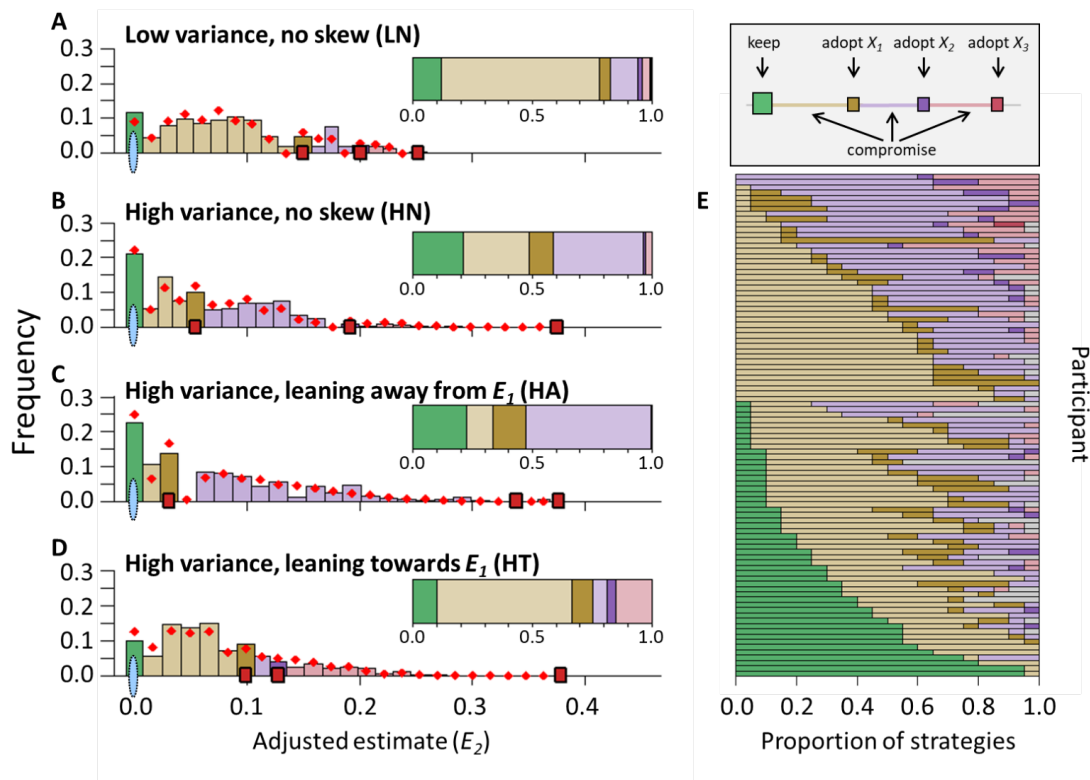


Figure 3.2. Adjustment strategies across conditions and participants. (A-D) Bars indicate the observed distribution of adjustments in individual rounds, expressed as the fraction of a participant’s first estimate (i.e., $|E_2 - E_1|/E_1$), per condition. The relative positions of the peer estimates (red squares) slightly varied across rounds (shown are their mean positions; see Methods for details). Bar and inset colors indicate the three strategies (i.e., keep, adopt, and compromise). Insets show the proportion of strategies per condition. Red diamonds show the predictions from the best cognitive model (see Cognitive Model), closely tracking the observed distributions. (E) The proportion of adjustment strategies per participant across all conditions. Participants (i.e., rows) are sorted according to their frequency of using the respective strategies.

assign more weight to estimates closer to each other (Fig. B4). This indicates that the observed deviations from the arithmetic mean in the experimental conditions (Fig. 3.1F; Fig. 3.2A-D) are not due to an inability to integrate multiple pieces of information. Rather, the stark differences between the experimental and control conditions suggest that people down-weight social information that is more distant from their first estimate, an effect known as ‘egocentric discounting’ (Yaniv, 1997; Yaniv and Kleinberger, 2000; Budescu et al., 2003; Harries et al., 2004; Yaniv, 2004; Yaniv and Milyavsky, 2007; Mannes, 2009; Moussaïd et al., 2013; Schultze et al., 2015; Jayles et al., 2017).

Cognitive Model. To investigate potential cognitive mechanisms underlying individuals’ integration of disparate social information, we developed a set of cognitive models unifying heuristic (i.e., keeping and adopting) with more complex strategies (i.e., compromising; 3.3A). Based on

our behavioral findings and previous literature, we identified four plausible assumptions. First, an individual's probability to keep its first estimate is higher, the closer the nearest peer (i.e., shorter distance between first and nearest peer estimate; Yaniv and Milyavsky, 2007; Schultze et al., 2015). Second, an individual's probability to adopt the estimate of the nearest peer is higher the closer the peer. Third, when compromising, peer estimates closer to an individual's first estimate are weighted more, capturing a possible 'confirmation' effect (Yaniv and Kleinberger, 2000; Yaniv, 2004; Moussaïd et al., 2013; Schultze et al., 2015; Jayles et al., 2017). Fourth, peer estimates closer to each other are weighted more, capturing a possible 'proximity' effect (Fig. B4; Yaniv and Milyavsky, 2007; Schultze et al., 2015). We formalize - and test - these assumptions below.

We cast these four assumptions into a Bayesian mixture model. First, individuals choose one of the three strategies with mixture probabilities $P(keep)$, $P(adopt)$, and $P(compromise)$ based on logistic functions (Rieskamp et al., 2003). Formally, the probability to keep is determined by: $P(keep) = S(\alpha_{keep} + \beta_{keep} \times d_1)$, where α_{keep} and β_{keep} are the intercept and slope of the standard logistic function (S), and d_1 is the absolute distance between the first estimate and the nearest peer estimate: $d_1 = |E_1 - X_1|$. Likewise, the probability to adopt the estimate of the nearest peer is also described by a logistic function of this distance: $P(adopt) = S(\alpha_{adopt} + \beta_{adopt} \times d_1)$. Note that we did not include the possibility to adopt the estimates of the peers that were not the closest, as this rarely happened (Fig. 3.2) and would make the model overly complicated.

The probability to compromise is now given by: $1 - P(keep) - P(adopt)$. Compromising is modelled as a Bayesian updating process (Fig. 3.3A; Toelch and Dolan, 2015; Park et al., 2017; Adjodah et al., 2017; Bang and Frith, 2017): individuals weigh their first estimate (E_1) and the peer estimates (X_{1-3}) and generate an updated (i.e., posterior) belief (Fig. 3.3A). We assume that an individual's initial (i.e., prior) belief E_p about the number of animals (N) follows a discretized normal distribution centered around the first estimate E_1 . The uncertainty of the belief is captured in the variance of the distribution (σ_p^2): $p(E_p|N) \sim Norm(E_1, \sigma_p^2)$.

Likewise, we model social information (SI_s) as discretized normal distributions around peer estimates (X_s), with uncertainty σ_s^2 : $p(SI_s|N) \sim Norm(X_s, \sigma_s^2)$, where the subscript s indexes each peer estimate. Both σ_p^2 and σ_s^2 are free parameters indicating how much weight individuals assign to the estimate, with high values (i.e., more uncertainty) indicating less weight. When $\sigma_p^2 < \sigma_s^2$, participants assign more weight to their own first estimate than to those of others.

Individuals further weigh each peer estimate based on two properties: its distance to their first

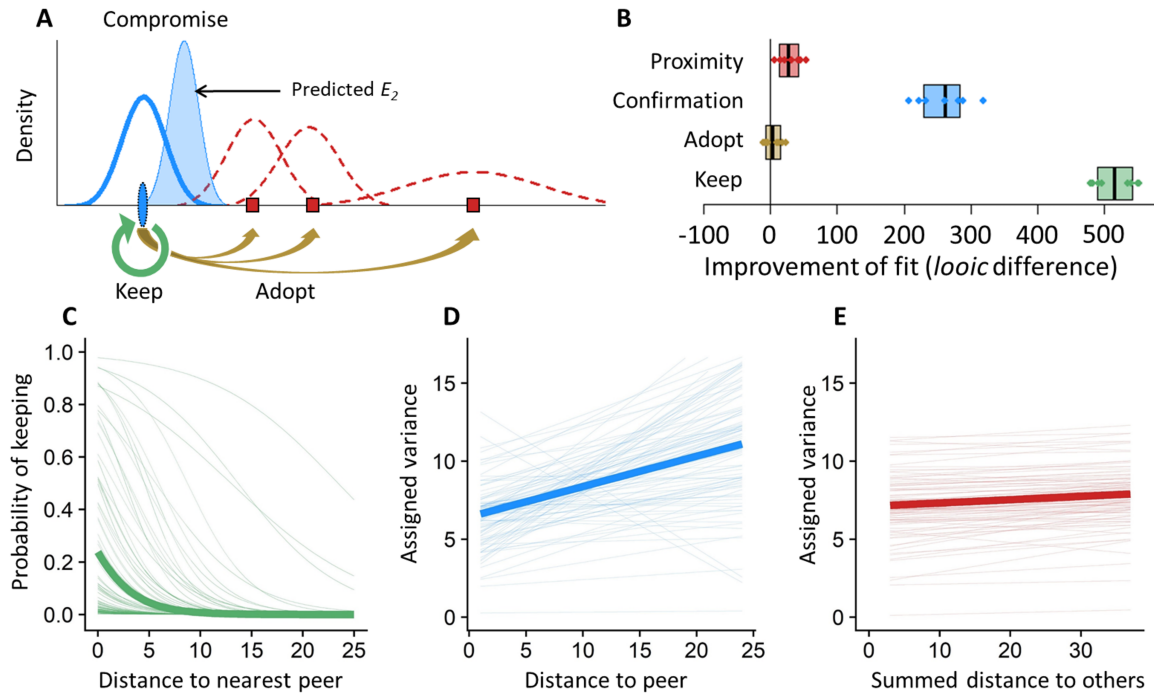


Figure 3.3. Cognitive model of the integration of disparate social information. (A) We model participants' estimate adjustments by combining heuristic updating strategies of keeping and adopting with more complex compromising strategies. The probability of keeping the first estimate (green arrow) was modelled as a function of the distance between a participants first estimate (E_1 ; blue oval) and the nearest peer (X_1). Similarly, the probability of adopting the nearest peer estimate (brown arrow) is a function of its distance to E_1 . Compromising entails taking a weighted average between E_1 and the peer estimates (X_i ; red squares): a participant's second estimate (E_2 ; transparent blue) is obtained by Bayesian updating using weighted means of E_1 and each X_i . Personal and social information are represented as probability density distributions with means at the observed estimates, and variances inversely related to their weight. Relative weights of social information depend on their distance to E_1 (degree of agreement with the participant; 'confirmation') and mean distance to other social information (i.e., degree of agreement with others; 'proximity'). (B) Improvement of model fit for each of the four features based on *looic* differences (Methods). Dots show the fit in improvement for pairs of models excluding and including each feature, and the box plots show the median improvement and IQR. (C-E) The main effect of each feature in the best-fitting model. (C) Fitted probability of keeping one's first estimate as a function of its distance to the nearest peer. (D and E) Variance assigned to peer estimates as a function of their distance to E_1 (D) and as their mean distance to other peers (E). Thin lines represent estimates for individuals and thick lines show group-level means.

estimate ('confirmation'), and its summed distance to the other two peer estimates ('proximity'). Confirmation-based weighting depends on the absolute distance between a peer estimate and an individual's first estimate ($d_s = |E_1 - X_s|$), $\sigma_s^2 = \alpha_s + \beta_{confirmation} \times d_s$, where α_s the intercept and $\beta_{confirmation}$ the slope of the distance weighting function. Likewise, proximity-based weighting of a peer estimate depends on its summed absolute distance to the other two peer estimates (e.g., $\tau_1 = |X_1 - X_2| + |X_1 - X_3|$), $\sigma_s^2 = \alpha_s + \beta_{proximity} \times \tau_s$. Both effects can also simultaneously shape uncertainty σ_s^2 in an additive fashion. Positive values of $\beta_{confirmation}$ or $\beta_{proximity}$ that the

weight assigned to peer estimates decreases as they become less consistent with an individual’s first estimate (i.e., confirmation) or with the other two peer estimates (i.e., proximity), respectively. This discounting would be, in turn, reflected in higher values of σ_s^2 . Bayes’ rule now allows to form an updated belief about the number of animals by integrating E_p and SI_s (see Supplementary Methods). Note that compromising can also result in instances of keeping or adopting.

We fit a series of models to our experimental data to estimate the parameter values of the different strategies: keeping, adopting, and Bayesian updating, including the confirmation and proximity effects. To account for individual differences in strategy use (Fig. 3.2E), we implement hierarchical Bayesian models (see Methods for implementation details). We evaluate the importance of the four model features ‘keep’, ‘adopt’, ‘confirmation’, and ‘proximity’ by calculating the leave-one-out cross-validation (looic; Vehtari et al., 2019) of the 16 models comprising all possible combinations of these features (Table B3).

Comparing the model fits in- and excluding each feature reveals that all features, except ‘adopt’, improve the model fit (Fig. 3.3B). Hence, the best-fitting model includes keeping and compromising based on ‘confirmation’ and ‘proximity’, but not adopting (Table B3). Overall, we find that participants (i) weight their own information more than the information of others, (ii) weigh social information more if it confirms their own information, and (iii) weigh social information more if it shows high inter-peer agreement (Table B4). Importantly, the best-fitting model closely predicts the mean adjustment across conditions (Fig. 3.1F; red vertical lines) as well as the distributions of adjustments in rounds across conditions (Fig. 3.2A-D; red diamonds). Moreover, it recovers the high inter-individual differences in mean adjustment and keeping probability (Fig. B5).

Figure 3.3C-E shows the effects of the three features of the best model. The highest improvement in model fit was found for keeping (Fig. 3.3B,C). Participants frequently kept their first estimate when they observed a peer in close agreement, but this likelihood sharply dropped with increasing distance to the nearest peer. When participants compromised, they assigned more weight to peers who more strongly agreed with them (Fig. 3.3D; note that the variance assigned to a piece of social information is inversely related to its weight). Both effects are indicative of a confirmation bias (i.e., favouring information that affirms one’s beliefs). Participants also assigned more weight to peers who showed a high agreement with other peers (a similar effect was found in the control trials without any personal information, Fig. B4). This ‘proximity’ effect was, however, much smaller than the effect ‘confirmation’ (Fig. 3.3B). For each of the model features, we observe

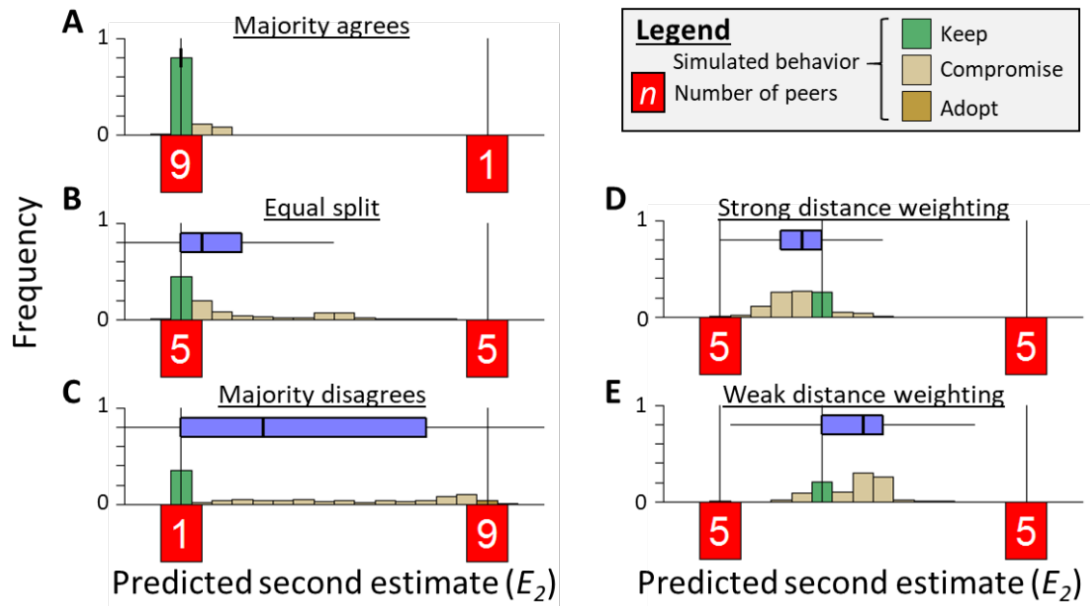


Figure 3.4. Simulated adjustments of agents observing ten pieces of social information in varying compositions. In the simulations, peer estimates were always Low or High (vertical black lines; the numbers in the red boxes correspond to the number of peers in each category; see Methods for details). In (A-C) The agent’s first estimate is also Low. In these situations, agents are very likely to keep their first estimate when they observe (A) a large majority agreeing with them, (B) half of the peers agreeing with them, and (C) only a small minority agreeing with them, though in the latter case, agents shifted substantially more. In (D-E) the agent’s first estimate is between Low and High (but closer to Low), and peers are equally split. (D) Agents with a strong confirmation bias (sampled from the upper 50% of the distribution; cf. steepest slopes in Fig. 3.3D) tend to adjust towards local extremes. (E) Agents with a weak confirmation bias (lower 50% of the distribution) tend to adjust towards the global mean estimate. In all panels, bars show distributions of predicted adjustments across 1,000 simulations. Boxplots summarize these distribution showing the median, IQR and the 1.5 IQR (whiskers).

substantial individual differences (indicated by the thin lines in Fig 3.3C-E). This is consistent with the high inter-individual differences in mean adjustment (Fig. B5). Finally, we note that the absence of an effect of adopting the closest peer estimate (Fig. 3.3B), is because these situations can be accounted for by adjustment through compromising.

These results indicate that our best-fitting model can account for the main patterns observed in our experimental conditions. Figure B5 shows that the model can also accurately predict participants’ mean adjustment and keep probability in the ‘filler’ rounds (where peer estimates were randomly selected from the prerecorded pool and frequently bracketed the participant’s first estimate). Figure B6 shows that the model can also recover a commonly-observed phenomena in estimation tasks, namely that mean adjustments are highest when social information is at intermediate distance from first estimates (Moussaïd et al., 2013; Jayles et al., 2017). Taken together, these results suggest that our model can be generalized to qualitatively different cases.

Simulations. The identified strategies of social information use allow us to predict how they may shape belief shifts in settings where people encounter peers with various levels of like-mindedness. In the following, we simulate the social information use in exemplary cases to explore what features of social information use and the social environment foster consensus, or, alternatively, lead to polarization. We simulate agents who, as in the experiment, observe other estimates and adjust their initial estimate. Agents observe ten pieces of social information, across five qualitative different settings: (i) a large majority agreeing with the agent, (ii) half of the peers agreeing with the agent, or (iii) only a small minority agreeing with the agent. Furthermore, we simulate the effect of a confirmation bias when the individual is slightly leaning towards one of two strongly disagreeing groups where we compare the adjustment of individuals with (iv) strong or (v) weak confirmation bias. In each setting, we simulated 1,000 agents whose adjustment strategies (i.e., their parameter setting) were sampled from the parent distributions of the best-fitting model (see Methods for details).

Figure 3.4 shows the predicted adjustments for the five scenarios. (i) When agents predominantly observe social information that agrees with their initial beliefs (i.e., mirroring a ‘filter bubble’; Pariser, 2011), they predominantly keep their initial estimate or, at most, adjust a tiny bit (Fig. 3.4A). (ii) Even when only half of the peers agree with them—and the other half strongly disagrees (a typical de-biasing attempt; Bozdag and van den Hoven, 2015)—agents only shift little, remaining far away from the global mean estimate (Fig. 3.4B). This suggests that even regular exposure to opposing information (e.g., from outside one’s filter bubble) does not result in strong behavioral adjustment. (iii) Even when only a small minority agrees with them, agents still have a high likelihood to keep their initial estimate, though overall adjustments towards the majority become more substantial (Fig. 3.4C). These results show the prominent role of the confirmation bias: confirming social information reinforces participants’ beliefs and prompts them to retain their beliefs, even if these beliefs reflect minority views.

To analyse the implications of a confirmation bias, we simulated adjustments for agents with high and low levels of this bias (i.e., individuals with steep and shallow slopes in Fig. 3.3D). The agents were located between two, equally-split, clusters of peers, but slightly closer to one of the clusters. Whereas agents with a strong confirmation bias adjusted their opinion towards the local cluster—and away from the global mean (Fig. 3.4D)—, agents with a weak confirmation bias adjusted towards the global mean estimate (Fig. 3.4E). These results suggest that a strong

confirmation bias can drive people to more extreme views, and, over time, increase polarization.

Discussion

This paper makes three novel contributions. First, we show that the impact of multiple sources of social information strongly depends on its distribution. Higher variance in social information reduces the participants' adjustments, while skewness decreases adjustments if a peer confirms their first estimate. Second, our cognitive models provide a unified and mechanistic account for how people integrate disparate social information, mapping out the determinants of adjustment strategies relying on simple heuristics (keeping first estimates and adopting those of others) and compromising (Bayesian inference). These models allow us to investigate participants' behaviour beyond accounting for mean estimate adjustments, and to capture how the weight of social information is independently shaped by its distance from people's initial beliefs and its proximity to other pieces of social information. Third and finally, our simulations illustrate how prioritizing confirmatory social information may exacerbate filter bubble effects and even cause individuals to adjust towards more extreme beliefs.

In real life, judging the accuracy of social information is often hard due to a lack of direct information on peer performance (or proxies thereof, like peer confidence, expertise or social status; Kurvers et al., 2019). In such cases, prioritizing social information that is consistent with other social information can be beneficial: when peers tend to make valid decisions, agreement reliably signals accuracy (Ravazzolo and Røisland, 2011; Mercier and Morin, 2019). It seems somewhat more puzzling why people show an even stronger tendency to prioritize their initial beliefs and social information confirming those beliefs, in a task where social information consists of the judgement of others who are incentivised to solve the same problem accurately. Indeed, from an external point of view, there is no reason to assume that one's initial judgment would be more accurate than those of others. People's overweighting of their initial beliefs might partly reflect a primacy effect (Asch, 1946); in our task, participants first formed their own beliefs before they observed social information. Another factor might be lacking access to the other's reasons for holding their opinions, as social information merely reflected the end-product of another person's decision process, whose identity and competence are unknown (Yaniv and Kleinberger, 2000). These processes may lead participants' to be (over) confident of their first judgments, and downweight social information especially when it is very distinct from one's own beliefs. Assessing the relative impact of confirmatory and disconfirming social cues on individuals' confidence in their

own judgement- as well as the order in which individuals access such information - would be an interesting topic for future research.

Our cognitive models provide a framework that unifies heuristic strategies (keeping and adopting) and strategies relying on weighted averaging between personal information and social information. Model selection revealed that assuming a separate heuristic for ‘adopting’ is not necessary, as Bayesian updating can already explain the instances of copying peer estimates in our experiment. This finding highlights to investigate whether identified qualitative different cases (here adopting), could also be explained by other established strategies (here compromising). Furthermore, cross-validation of our models by predicting behaviour in the ‘filler rounds’ underscores their ability to go beyond mere description, and suggest they can be generalized to settings beyond our experimental conditions (Fig. B5; see also Fig. B6). Moreover, our hierarchical modeling approach allowed for accounting for individual differences in social information use, a regularly observed but poorly understood phenomenon (Efferson et al., 2008; Molleman et al., 2014; Mesoudi et al., 2016). In sum, our models provide a detailed and predictive account of people’s adjustment strategies, contributing to understanding the hitherto understudied computational and cognitive mechanisms underlying social information use (Heyes, 2016).

Our simulations predict that disparate social information changes people’s minds only to a limited degree, even when it signals that they hold minority views (Fig. 3.4). Moreover, under certain conditions, observing balanced social information can in principle lead individuals to take more extreme views (Fig. 3.4D,E). These findings suggest that efforts to de-bias online information that present people with balanced views (Bozdag and van den Hoven, 2015) might not suffice to break filter bubble effects and dynamics of polarization. Our simulations suggest that people’s minds are most likely changed in social environments where none of their peers agrees with them. Future empirical work should test these predictions, as well as the extent to which our results generalise to other (possibly non-WEIRD; Henrich et al., 2010) populations and, perhaps more importantly, other - and possibly richer - domains of decision making (e.g. matters of taste, choosing a product to buy, or moral, emotive or political issues; Analytis et al., 2018). It seems plausible that in many important real-world contexts, the integration of disparate social information can be further hampered due to ‘motivated reasoning’ (selectively bringing in reasons countering information that conflicts with one’s worldview; Bail et al 2018) or when observed peers belong to an out-group (Votruba and Kwan, 2015; Ostrom et al., 1993; Guilbeault et al., 2018). On the other hand, at least some people might be susceptible to incorporating disparate views, when these views are

accompanied by good arguments (Laughlin, 2011; Mercier et al., 2017).

Methods

Behavioral experiment. We recruited 100 participants from Amazon Mechanical Turk (MTurk), restricted to US citizens. By clicking the link to the experimental pages, participants confirmed informed consent. Five participants dropped out during the task and did not receive any payment, resulting in 95 participants (57% male, mean age: 35.8; s.d.=10.7). The online experiment was programmed in LIONESS Lab (Giamattei et al., 2019). Ethical approval was obtained from the Institutional Review Board of the Max Planck Institute for Human Development Berlin (ARC 2017/18).

The experimental task is based on a validated perceptual judgment paradigm for quantifying social information use (Fig. 3.1A-C). In each of 30 rounds (5 per experimental condition, plus 10 ‘filler’ rounds with random social information, see below), participants observed an image with 50-100 animals for 6 seconds and had to estimate how many animals there were. They entered their estimates with a slider limited from 1 to 150. After submitting their first estimate (E_1), they observed the estimates of three other participants (X_1, X_2, X_3) who had completed the same task before but without receiving social information. After observing the social information, participants provided a second estimate (E_2). Participants were rewarded for accuracy, earning 100 points if their estimate was exactly correct (both for E_1 and E_2). For each animal they were off, five points were subtracted (but earnings in a round could not drop below zero). At the end of the session, one decision was randomly chosen from each of the experimental ‘blocks’ (see Supplementary Methods) for bonus payment (100 points = \$1.00), which came on top of on top of a flat fee of \$4.50. Total earnings ranged from \$4.50 to \$7.00 (average \$5.50). Participants took, on average, 35 minutes, resulting in an hourly wage of \$9.50. Experimental sessions ended with a short questionnaire in which we recorded participants’ age and gender, and measured individualism (Triandis and Gelfand, 1998), social conformity (Mehrabian and Stefl, 1995), and resistance to peer influence (Fig. B1; Steinberg and Monahan, 2007).

We used four experimental conditions (Fig. 3.1E) systematically varying the variance and skewness of the distributions across conditions, while keeping the distance between the mean social information (X) and E_1 constant. To achieve experimental control without deception, we first recorded a large ($N=100$) pool of estimates made by MTurkers for each image shown in the main experiment. In a given trial, the three pieces of social information were selected based on the

participant’s first estimate and the experimental condition. Based on these, the three estimates that most closely matched the conditions were selected (for implementation details, see Supplementary Methods). This procedure resulted in clearly defined experimental conditions (Fig. B7). We randomly shuffled the order of experimental conditions across rounds and held this order fixed for all participants.

Cognitive Model. We analysed the cognitive model with a hierarchical Bayesian MCMC technique implemented with “RStan” in R (R Core Team, 2019; Stan Development Team, 2018, for implementation of the sampler and the hierarchical model structure, see Supplementary Method). We investigated the predictive power of the four model features ‘keep’, ‘adopt’, ‘confirmation’ and ‘proximity’ by calculating the *loaic* (Vehtari et al., 2019) of the 16 model variants comprising all possible combinations of these features (Table B3). We quantified the importance of a feature by calculating the average reduction of the *loaic* when the feature was included (Fig. 3.2B). We report the fittings of the model with the lowest *loaic* (see Supplementary Methods).

We generated predictions of the best fitting model (i.e., lowest *loaic*) by calculating the probability density function for each participant and round. This density function was based on the mean posterior parameter estimates for each participant (see Table B4 for characteristics of the parameter parent distributions), the first estimate (E_1) and the observed social information (X_s). To account for stochasticity, the model predictions in Figure 3.1F and Figure 3.2C (red dots) are based on 10 samples of estimates from each density function. To analyse the prediction of the best-fitting model for the experimental conditions versus ‘filler’ rounds we once sampled from each density function of each round and participant and calculated the actual and predicted individual-level mean adjustment and keep proportions (Fig. B4).

Simulations. In each of the settings, individuals were endowed with adjustment strategies whose parameter values were sampled from the parent distributions from the best-fitting model (Table B4), assuming no correlations between the parameters. For each agent we emulated a first estimate, and simulated their adjustment given the social environment. For Fig. 3.4, the social environment consisted of ten pieces of social information, either agreeing (i.e., an estimate of 50) or disagreeing (i.e., an estimate of 65) with the focal agent. For Fig. 3.4D and 3.4E we sampled the distance weighting parameter from the higher and lower half of the parent distribution, respectively. The individuals’ first estimate in these settings was 55.

References

- Adjodah, D., Leng, Y., Chong, S. K., Krafft, P., and Pentland, A. (2017). Social bayesian learning in the wisdom of the crowd.
- Aitchison, L., Bang, D., Bahrami, B., and Latham, P. E. (2015). Doubly bayesian analysis of confidence in perceptual decision-making. *PLoS Computational Biology*, 11(10):e1004519.
- Analytis, P. P., Barkoczi, D., and Herzog, S. M. (2018). Social learning strategies for matters of taste. *Nature Human Behaviour*, 2(6):415–428.
- Aplin, L. M., Farine, D. R., Morand-Ferron, J., Cockburn, A., Thornton, A., and Sheldon, B. C. (2015). Experimentally induced innovations lead to persistent culture via conformity in wild birds. *Nature*, 518(7540):538–541.
- Asch, S. E. (1946). Forming impressions of personality. *The Journal of Abnormal and Social Psychology*, 41(3):258–290.
- Asch, S. E. (1956). Studies of independence and conformity: I. a minority of one against a unanimous majority. *Psychological Monographs: General and Applied*, 70(9):1–70.
- Bahrami, B., Olsen, K., Latham, P. E., Roepstorff, A., Rees, G., and Frith, C. D. (2010). Optimally Interacting Minds. *Science*, 329(5995):1081–1085. WOS:000281253500042.
- Bail, C. A., Argyle, L. P., Brown, T. W., Bumpus, J. P., Chen, H., Hunzaker, M. F., Lee, J., Mann, M., Merhout, F., and Volfovsky, A. (2018). Exposure to opposing views on social media can increase political polarization. *Proceedings of the National Academy of Sciences*, 115(37):9216–9221.
- Bakshy, E., Rosenn, I., Marlow, C., and Adamic, L. (2012). The role of social networks in information diffusion. In *Proceedings of the 21st International Conference on World Wide Web*, pages 519–528, New York, NY, USA. ACM, ACM.
- Bang, D. and Frith, C. D. (2017). Making better decisions in groups. *Royal Society Open Science*, 4(8):170193.
- Bednarik, P. and Schultze, T. (2015). The effectiveness of imperfect weighting in advice taking. *Judgment and Decision Making*, 10(3):265–276.
- Bloomfield, R. and Hales, J. (2009). An experimental investigation of the positive and negative effects of mutual observation. *The Accounting Review*, 84(2):331–354.
- Bonaccio, S. and 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.
- Bond, R. (2005). Group size and conformity. *Group Processes & Intergroup Relations*, 8(4):331–354.
- Boyd, R. and Richerson, P. J. (1985). *Culture and the evolutionary process*. University of Chicago Press.
- Boyd, R., Richerson, P. J., and Henrich, J. (2011). The cultural niche: Why social learning is essential for human adaptation. *Proceedings of the National Academy of Sciences*, 108:10918.
- Bozdag, E. and van den Hoven, J. (2015). Breaking the filter bubble: democracy and design. *Ethics and Information Technology*, 17(4):249–265.
- Budescu, D. V., Rantilla, A. K., Yu, H.-T., and Karelitz, T. M. (2003). The effects of asymmetry among advisors on the aggregation of their opinions. *Organizational Behavior and Human Decision Processes*, 90(1):178–194.
- Budescu, D. V. and Yu, H.-T. (2007). Aggregation of opinions based on correlated cues and advisors. *Journal of Behavioral Decision Making*, 20(2):153–177.
- Danchin, E., Nöbel, S., Pocheville, A., Dagaëff, A.-C., Demay, L., Alphand, M., Ranty-Roby, S., van Renssen, L., Monier, M., Gazagne, E., et al. (2018). Cultural flies: Conformist social learning in fruitflies predicts long-lasting mate-choice traditions. *Science*, 362(6418):1025–1030.
- Derex, M., Bonnefon, J.-F., Boyd, R., and Mesoudi, A. (2019). Causal understanding is not necessary for the improvement of culturally evolving technology. *Nature Human Behaviour*, 3(5):446–452.
- Deryugina, T. and Shurchkov, O. (2016). The effect of information provision on public consensus about climate change. *PLoS ONE*, 11(4):e0151469.
- Devenow, A. and Welch, I. (1996). Rational herding in financial economics. *European Economic Review*, 40(3):603 – 615. Papers and Proceedings of the Tenth Annual Congress of the European Economic Association.
- Efferson, C., Lalive, R., and Fehr, E. (2008). The coevolution of cultural groups and ingroup favoritism.

- Science*, 321(5897):1844–1849.
- Efferson, C., Vogt, S., and Fehr, E. (2019). The promise and the peril of using social influence to reverse harmful traditions. *Nature Human Behaviour*, pages 1–14.
- Giamattei, M., Molleman, L., Seyed Yahosseini, K., and Gächter, S. (2019). Lioness lab—a free web-based platform for conducting interactive experiments online.
- Guilbeault, D., Becker, J., and Centola, D. (2018). Social learning and partisan bias in the interpretation of climate trends. *Proceedings of the National Academy of Sciences*, 115(39):9714–9719.
- Harries, C., Yaniv, I., and Harvey, N. (2004). Combining advice: The weight of a dissenting opinion in the consensus. *Journal of Behavioral Decision Making*, 17(5):333–348.
- Henrich, J., Heine, S. J., and Norenzayan, A. (2010). Most people are not weird. *Nature*, 466(7302):29.
- Heyes, C. (2016). Blackboxing: social learning strategies and cultural evolution. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 371(1693):20150369.
- Hoppitt, W. and Laland, K. N. (2013). *Social learning: an introduction to mechanisms, methods, and models*. Princeton University Press.
- Jayles, B., Kim, H.-r., Escobedo, R., Cezera, S., Blanchet, A., Kameda, T., Sire, C., and Theraulaz, G. (2017). How social information can improve estimation accuracy in human groups. *Proceedings of the National Academy of Sciences*, 114(47):12620–12625.
- Kendal, R. L., Boogert, N. J., Rendell, L., Laland, K. N., Webster, M., and Jones, P. L. (2018). Social Learning Strategies: Bridge-Building between Fields. *Trends in Cognitive Sciences*, 22(7):651–665.
- Kerr, J. R. and Wilson, M. S. (2018). Perceptions of scientific consensus do not predict later beliefs about the reality of climate change: A test of the gateway belief model using cross-lagged panel analysis. *Journal of Environmental Psychology*, 59:107–110.
- Kurvers, R. H., Wolf, M., and Krause, J. (2014). Humans use social information to adjust their quorum thresholds adaptively in a simulated predator detection experiment. *Behavioral Ecology and Sociobiology*, 68(3):449–456.
- Kurvers, R. H. J. M., Herzog, S. M., Hertwig, R., Krause, J., Moussaid, M., Argenziano, G., Zalaudek, I., Carney, P. A., and Wolf, M. (2019). How to detect high-performing individuals and groups: Decision similarity predicts accuracy. *Science Advances*, 5(11).
- Larrick, R. P. and Soll, J. B. (2006). Intuitions about combining opinions: Misappreciation of the averaging principle. *Management Science*, 52(1):111–127.
- Laughlin, P. R. (2011). *Group problem solving*. Princeton University Press, Princeton.
- Lewandowsky, S., Gignac, G. E., and Vaughan, S. (2013). The pivotal role of perceived scientific consensus in acceptance of science. *Nature Climate Change*, 3(4):399–404.
- 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.
- Mehrabian, A. and Steff, C. A. (1995). Basic temperament components of loneliness, shyness, and conformity. *Social Behavior and Personality: an international journal*, 23(3):253–263.
- Mercier, H., Dezechache, G., and Scott-Phillips, T. (2017). Strategically communicating minds. *Current Directions in Psychological Science*, 26(5):411–416.
- Mercier, H. and Morin, O. (2019). Majority rules: how good are we at aggregating convergent opinions? *Evolutionary Human Sciences*, 1.
- Mesoudi, A. (2011). An experimental comparison of human social learning strategies: payoff-biased social learning is adaptive but underused. *Evolution and Human Behavior*, 32(5):334–342.
- Mesoudi, A., Chang, L., Dall, S. R., and Thornton, A. (2016). The evolution of individual and cultural variation in social learning. *Trends in Ecology & Evolution*, 31(3):215–225.
- Molleman, L., Kurvers, R. H., and van den Bos, W. (2019). Unleashing the beast: a brief measure of human social information use. *Evolution and Human Behavior*, 40(5):492 – 499.
- Molleman, L., Van den Berg, P., and Weissing, F. J. (2014). Consistent individual differences in human social learning strategies. *Nature Communications*, 5(3570).
- Moussaïd, M., Herzog, S. M., Kämmer, J. E., and Hertwig, R. (2017). Reach and speed of judgment propagation in the laboratory. *Proceedings of the National Academy of Sciences*.
- Moussaïd, M., Kämmer, J. E., Analytis, P. P., and Neth, H. (2013). Social Influence and the Collective Dynamics of Opinion Formation. *PLoS ONE*, 8(11):e78433.
- Ostrom, T. M., Carpenter, S. L., Sedikides, C., and Li, F. (1993). Differential processing of in-group and

- out-group information. *Journal of Personality and Social Psychology*, 64(1):21–34.
- Page, S. E. (2008). *The Difference: How the Power of Diversity Creates Better Groups, Firms, Schools, and Societies-New Edition*. Princeton University Press.
- Pariser, E. (2011). *The filter bubble: What the Internet is hiding from you*. Penguin Press, New York, NY.
- Park, S. A., Gojame, S., O'Connor, D. A., and Dreher, J.-C. (2017). Integration of individual and social information for decision-making in groups of different sizes. *PLoS Biology*, 15(6):e2001958.
- R Core Team (2019). *R: A language and environment for statistical computing*. R Foundation for Statistical Computing, Vienna, Austria.
- Ravazzolo, F. and Røisland, Ø. (2011). Why do people place lower weight on advice far from their own initial opinion? *Economics Letters*, 112(1):63–66.
- Rendell, L., Boyd, R., Cownden, D., Enquist, M., Eriksson, K., Feldman, M. W., Fogarty, L., Ghirlanda, S., Lillicrap, T., and Laland, K. N. (2010). Why Copy Others? Insights from the Social Learning Strategies Tournament. *Science*, 328(5975):208.
- Rieskamp, J., Busemeyer, J. R., and Laine, T. (2003). How Do People Learn to Allocate Resources? Comparing Two Learning Theories. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 29(6):1066–1081.
- Schultze, T., Rakotoarisoa, A.-F., and Schulz-Hardt, S. (2015). Effects of distance between initial estimates and advice on advice utilization. *Judgment and Decision Making*, 10(2):144–171.
- Shea, N., Boldt, A., Bang, D., Yeung, N., Heyes, C., and Frith, C. D. (2014). Supra-personal cognitive control and metacognition. *Trends in Cognitive Sciences*, 18(4):186–193.
- Soll, J. B. and Larrick, R. P. (2009). Strategies for revising judgment: How (and how well) people use others' opinions. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 35(3):780–805.
- Stan Development Team (2018). RStan: the R interface to Stan. R package version 2.18.2.
- Steinberg, L. and Monahan, K. C. (2007). Age differences in resistance to peer influence. *Developmental Psychology*, 43(6):1531–1543.
- Stewart, A. J., Mosleh, M., Diakonova, M., Arechar, A. A., Rand, D. G., and Plotkin, J. B. (2019). Information gerrymandering and undemocratic decisions. *Nature*, 573(7772):117–121.
- Sunstein, C. R. (2007). *Republic. com 2.0*. Princeton University Press, Princeton.
- Surowiecki, J. (2004). *The Wisdom of Crowds: Why the Many Are Smarter Than the Few and How Collective Wisdom Shapes Business, Economies, Societies and Nations*. Doubleday, New York.
- Toelch, U. and Dolan, R. J. (2015). Informational and normative influences in conformity from a neurocomputational perspective. *Trends in cognitive sciences*, 19(10):579–589.
- Triandis, H. C. and Gelfand, M. J. (1998). Converging measurement of horizontal and vertical individualism and collectivism. *Journal of Personality and Social Psychology*, 74(1):118–128.
- Tump, A., Wolf, M., Krause, J., and Kurvers, R. H. (2018). Individuals fail to reap the collective benefits of diversity because of over-reliance on personal information. *Journal of the Royal Society Interface*, 15(142):20180155.
- van den Berg, P., Molleman, L., and Weissing, F. J. (2015). Focus on the success of others leads to selfish behavior. *Proceedings of the National Academy of Sciences*, 112(9):2912–2917.
- Vehtari, A., Gabry, J., Yao, Y., and Gelman, A. (2019). loo: Efficient leave-one-out cross-validation and waic for bayesian models. R package version 2.1.0.
- Votruba, A. M. and Kwan, V. S. (2015). Disagreeing on whether agreement is persuasive: Perceptions of expert group decisions. *PLoS ONE*, 10(3):e0121426.
- Yaniv, I. (1997). Weighting and trimming: Heuristics for aggregating judgments under uncertainty. *Organizational Behavior and Human Decision Processes*, 69(3):237–249.
- Yaniv, I. (2004). Receiving other people's advice: Influence and benefit. *Organizational Behavior and Human Decision Processes*, 93(1):1 – 13.
- Yaniv, I. and Kleinberger, E. (2000). Advice taking in decision making: Egocentric discounting and reputation formation. *Organizational Behavior and Human Decision Processes*, 83(2):260 – 281.
- Yaniv, I. and Milyavsky, M. (2007). Using advice from multiple sources to revise and improve judgments. *Organizational Behavior and Human Decision Processes*, 103(1):104–120.

Chapter 4

Wise or mad crowds? The cognitive mechanisms underlying information cascades

Tump, A.N., Pleskac, T., & Kurvers, R.H.J.M.

This chapter can be accessed at:

<https://doi.org/10.31234/osf.io/6vt2p>.

Abstract

Whether getting vaccinated, buying stocks, or crossing streets, people rarely make decisions alone. Rather, multiple people decide sequentially, setting the stage for information cascades whereby early-deciding individuals can influence others' choices. To understand how information cascades through social systems, it is essential to capture the dynamics of the decision making process. We introduce the social drift-diffusion model to capture these dynamics. We tested our model using a sequential choice task. The model was able to recover the dynamics of the social decision making process, accurately capturing how individuals integrate personal and social information dynamically over time and when they timed their decisions. Our results show the importance of the interrelationships between accuracy, confidence, and response time in shaping the quality of information cascades. The model reveals the importance of capturing the dynamics of decision processes to understand how information cascades in social systems, paving the way for applications in other social systems.

Acknowledgement: We thank Max Wolf for helpful feedback, Susannah Goss for editing the manuscript, Oliver Krüger for hosting the experiments, and all tutors of the 2017 'Basismodul Biologie' at the University of Bielefeld, Germany. R.H.J.M.K. acknowledges funding from the German Research Foundation (DFG, grant number: KU 3369/1-1).

Data and Code Availability: The data and code to implement all analysis can be accessed in

https://osf.io/ejfm4/?view_only=9748c5bf32874c6f9745e3af280ca227

Author Contributions: Conceptualization: A.N.T., R.H.J.M.K., & T.J.P.; Methodology: A.N.T., R.H.J.M.K., & T.J.P.; Software: A.N.T.; Data collection & curation: A.N.T.; Formal analysis: A.N.T., R.H.J.M.K., & T.J.P.; Writing - Original Draft A.N.T. and R.H.J.M.K.; Writing - Reviewing & Editing - A.N.T., R.H.J.M.K., & T.J.P.; Supervision: R.H.J.M.K.

Introduction

In many situations—be they financial investments, consumer choices, or simply crossing the street—one is generally not making a decision alone. Rather, there are multiple others present each making their own decisions. In such situations, decision makers can observe the choices of others and use that information to inform their own decisions. Early-deciding individuals can thereby trigger *information cascades*, in which later-deciding individuals adopt earlier choices, potentially creating a situation where, in the extreme case, everyone does what everyone else is doing, even at the expense of abandoning their private information (Banerjee, 1992; Anderson and Holt, 1997; Bikhchandani et al., 1998; Gallup et al., 2012).

Yet for a myriad of reasons—from limited time and computational resources to biases in the decision process—people’s choices do not always perfectly reflect the true state of the world. Information cascades can thus promote both positive and negative outcomes: in online environments, for example, both true and fake news can spread quickly (Vosoughi et al., 2018); in offline environments, the behaviour of initial pedestrians crossing a road can amplify both safe and risky behaviours in other pedestrians (Faria et al., 2010; Pfeffer and Hunter, 2013). Understanding the conditions leading to positive and negative information cascades is crucial across many domains, including financial markets (Welch, 2000; Shiller, 2002), consumer preferences (Chen, 2008), political opinion formation (Battaglini, 2005), and opinion dynamics in social networks (Xiong and Liu, 2014).

To understand the conditions underlying positive and negative information cascades, we need to comprehend the timing of individual decisions as well as how individuals integrate personal and social information (i.e., other people’s decisions) dynamically over time. We, currently, however, lack a detailed understanding of the individual decision process in sequential choice paradigms. Many models of information cascades assume a random decision order and are thus ill-equipped to predict who will respond earlier and why (e.g., Anderson and Holt, 1997; Banerjee, 1992; Bikhchandani et al., 1998; Deneubourg et al., 1990; Mann, 2018; Sumpter and Pratt, 2008). When models of information cascades do refer to the timing of decisions, they do so from an optimal Bayesian perspective based on the quality of each individual’s private information (e.g., Chamley and Gale, 1994; Gul and Lundholm, 1995; Zhang, 1997; Ziegelmeyer et al., 2005). Yet we know that people’s actual choice behaviour often deviates systematically from optimal Bayesian models (Hertwig et al., 2019; Pleskac and Busemeyer, 2010; Tversky and Kahneman, 1974).

To address these shortcomings, we developed a dynamic theory of social decision making by focusing on each individual's decision process. As a basis, we took a well-established modeling framework of individual decision making that models decisions as a dynamic process in which information is accumulated as evidence over time until a threshold is reached (e.g., Edwards, 1965; Laming, 1968; Link and Heath, 1975; Ratcliff, 1978; Stone, 1960; Usher and McClelland, 2001). This evidence accumulation process has been successful in accounting for a wide range of decisions in domains including perception (Ratcliff and Smith, 2004), memory (Ratcliff, 1978), categorization (Nosofsky and Palmeri, 1997), preference (Busemeyer and Diederich, 2002; Busemeyer and Townsend, 1993; Konovalov and Krajbich, 2019), inference (Pleskac and Busemeyer, 2010), and has successfully been applied to analyse the influence of static social information (Germar et al., 2014; Toelch et al., 2018). We extended this evidence accumulation framework by showing how the choices of others are integrated with personal information and together accumulated as evidence. This approach provides a process-level account of the choices and response times of individuals in dynamic social systems. We tested the model in an empirical study. Findings showed that participants self-organize based on the quality of their personal information so that later deciders benefit from observing the choices of early deciders. Fitting the model to the data allowed us to test several hypotheses about how individuals simultaneously combine personal and social information, and how they time their decision in groups. In addition, we reveal mechanisms leading to the amplification of correct or incorrect cascading information.

The Social Drift-Diffusion Model

Models of the evidence accumulation process during decision making include the drift-diffusion model (DDM; Ratcliff, 1978; Ratcliff and McKoon, 2008), the linear ballistic accumulator model (Brown and Heathcote, 2008), and the leaky competing accumulator model (Usher and McClelland, 2001). Most of these models can, in principle, be extended to model a social system. Here, we focus on the DDM, arguably the most successful framework for accounting for human choice behaviour, including some of the most basic aspects of the decision process, such as the speed-accuracy trade-off (Ratcliff and Smith, 2004; Voss et al., 2004), the construction of preferences (Busemeyer and Townsend, 1993), the formation of confidence judgements (Pleskac and Busemeyer, 2010), the emergence of response biases (Leite and Ratcliff, 2011; Pleskac et al., 2018), and how attention guides the evidence accumulation process (Diederich, 1997; Krajbich and Rangel, 2011).

According to the DDM, people faced with a choice between two options, *A* or *B*, base their

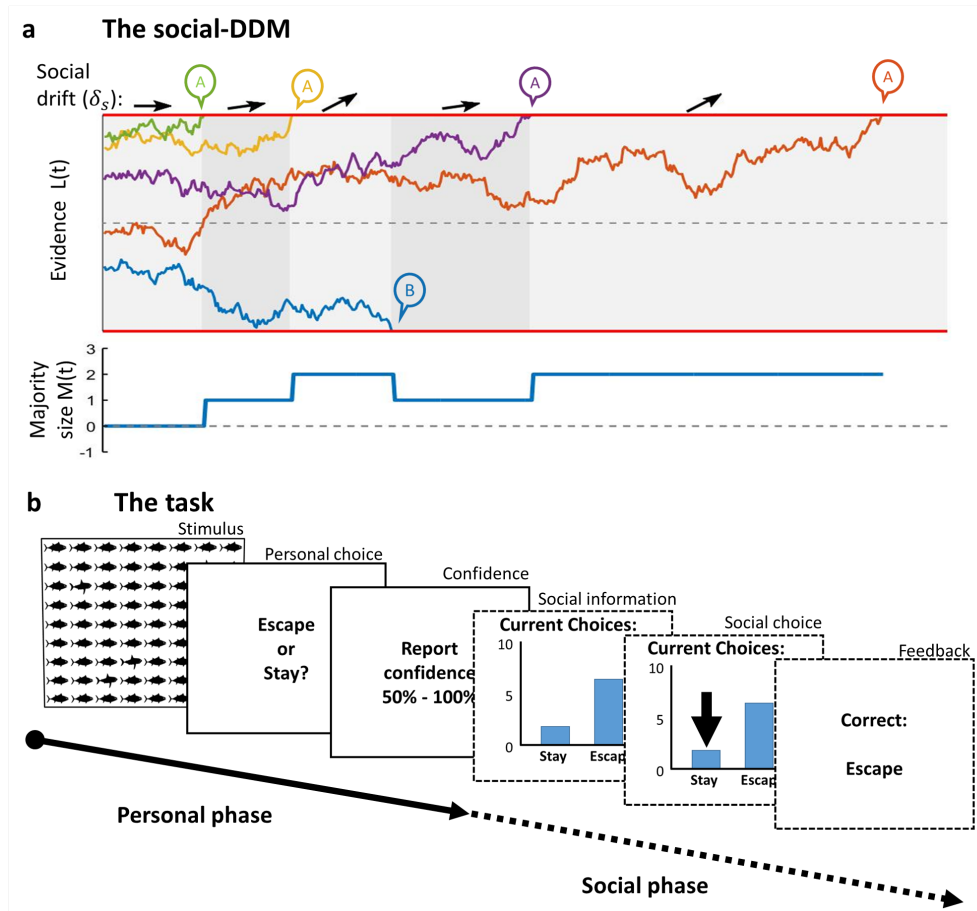


Figure 4.1. Illustration of the social DDM and the experimental paradigm. (a) A generic example of the social DDM with five individuals, each represented by a jagged line. The start point of each individual indicates the personal evidence accumulated up to that point. At the start, no individual exceeds the choice threshold and social information is absent, implying no social drift (as indicated by the horizontal arrow). Individuals who begin close to either of the thresholds (red lines) are likely to choose early, providing social information for undecided individuals. This social information impacts the rate of evidence accumulation, with the drift rate shifting towards the choice threshold favoured by the majority (as indicated by the arrow pointing upwards). (b) The stages of the predator detection task. During the personal phase, individuals briefly observe a grid of ‘sharks’ and ‘tuna.’ They then make a personal decision whether to ‘Stay’ or ‘Escape’ and report their confidence in that decision. In the subsequent social phase, they are asked to make a second decision on whether to ‘Stay’ or ‘Escape,’ but now they can freely time their decisions and simultaneously observe the choices of others before doing so. Finally, the correct answer is displayed, and the next trial begins (with 40 trials in total).

choice on an internal level of evidence. Initially, people can have a bias and lean towards either option. This is modeled as an initial level of evidence. Over time, people extract further information about the options and accumulate this information as evidence. This accumulation gives rise to an evolving (latent) level of evidence, as depicted by the jagged line in Figure 4.1a. The jaggedness arises because each sample of evidence is noisy (i.e., the stimuli itself and the cognitive and neural processes introduce variability into the evidence accumulation). Once a choice threshold

has been reached, a decision is made. If the accumulated evidence reaches the upper threshold, option A is selected; if it crosses the lower threshold, option B is selected. The time it takes for the evidence to reach either threshold is the predicted response time. In the social DDM, we modify this framework to cover multiple individuals accumulating evidence at the same time (Fig. 4.1a). In this case, the evidence comes from two sources: personal information, gathered from sampling the physical environment (e.g., for visual or auditory cues), and social information, gathered by observing the behaviour of others (Dall et al., 2005; Galef and Laland, 2005).

Formally, we denote the cumulative evidence at time point t as $L(t)$. At the start, individuals may favour one option over the other, described by their start point $L(t = 0) = \beta$. Here, their start point is based on previously collected personal information and is estimated from confidence ratings provided during the initial stage of the decision task. However, the start point can also represent initial biases towards either option (e.g., Voss et al., 2004). At each time step Δt , the current state of evidence $L(t)$ is updated by sampling new evidence until a decision is made (i.e., until the level of evidence reaches the choice threshold θ):

$$L(t + \Delta t) = L(t) + [\delta_p + \delta_s(t)] \times \Delta t + \sqrt{\Delta t} \times \epsilon, \quad (4.1)$$

where ϵ is Gaussian white noise (i.e., the diffusion process) with a mean of 0 and a variance of 1. The parameters δ_p and $\delta_s(t)$ correspond to the strength of the personal and social information uptake, respectively. Personal information uptake describes the integration of information extracted directly from the physical environment, as well as the evaluation of information from memory. Social information is defined as the size of the majority of individuals $M(t)$ who already decided at time point t (see also Bikhchandani et al., 1998):

$$M(t) = N_A(t) - N_B(t), \quad (4.2)$$

where $N_A(t)$ and $N_B(t)$ are the number of individuals who have already decided for option A or B, respectively. The impact of majority size on the social drift rate is described by a power function, analog to Latané (1981):

$$\delta_s(t) = s \times M(t)^q. \quad (4.3)$$

The parameter s is a scaling factor that influences the strength of the social drift; q governs the shape of the power function. When $q = 1$, each additional choice for the majority option has the same influence on the social drift rate (i.e., a linear effect); when $q > 1$ ($q < 1$), each additional choice for the majority option has an increasingly stronger (weaker) impact on the social drift

Table 4.1. Description of the parameters of the social DDM

Model feature	Parameter	Description
Nondecision time	τ	Response latency (e.g., motor response time). The parameter τ describes the time relative to the individual's fastest response.
Start point	$\beta = \frac{1}{1+e^{-a(C-b)}}$	The start point is a function of the confidence in the personal choice C , which ranges from highly confident but incorrect to highly confident and correct (Fig. 4.4b). The parameter a determines how sensitive the start point is to changes in confidence; b captures other factors besides confidence in the personal decision that impact the start point.
Personal drift rate	δ_p	The average rate of evidence accumulation supporting the personal choice (Fig. 4.4c).
Social drift rate	$\delta_s = s \times M(t)^q$	The social drift rate describes the impact of social information, with s being a scaling parameter that influences the strength of the social drift rate, and q being a parameter that shapes the power function describing the relationship of majority size $M(t)$ and social drift rate (Fig. 4.4d).
Choice threshold	θ	The amount of evidence an individual has to accumulate to make a decision; θ ($-\theta$) reflects the correct (incorrect) choice threshold (Fig. 4.4e).

rate. Note that, in contrast to the individual drift rate, the social drift rate can vary over time (indicated by the changing direction of the arrows in Fig. 4.1a). By incorporating a social drift into the classical DDM, the social DDM can account for individuals being emitters and receivers of social information and thereby capture the dynamic information exchange among group members.

In sum, the social DDM characterizes (i) how individuals incorporate personal information with the parameters β and δ_p , (ii) how individuals incorporate social information depending on the majority size via the parameters s and q , and (iii) individuals' willingness to wait for social information with the parameter θ (see Table 4.1 for all parameter descriptions).

The predator detection task

We tested the social DDM in an empirical study (see Fig. 4.1b; see Methods for full details). In brief, participants were divided into groups of varying sizes ('small', 'medium', or 'large'). Each group of participants was seated together in a single room, facing a large screen. Participants were asked to imagine being a fish in a school facing a choice between two alternatives—namely whether to escape or not—depending on the presence of predators, in this case, sharks. They were instructed to escape when five or more sharks were present and to stay when four or fewer sharks were present. At each trial, participants were shown—for 2 seconds—a grid with a varying number (3, 4, 6, or 7) of sharks hidden among harmless fish. Participants first made a personal choice on whether to 'Stay' or 'Escape' and then reported their confidence in that choice on a scale from 50% to 100%. They then entered the social phase, in which they had a maximum of 20 seconds to make a second decision on whether to 'Stay' or 'Escape', but without seeing the grid again. Instead, the display showed a count of the number of choices for each option. Participants were free to enter their choice at any point in time; they could thus respond early (thereby providing social information) or wait to observe the decisions of others. However, they could only decide once. Finally, we provided feedback on the correct choice.

Results

Empirical results: groups show beneficial self-organization according to information quality

Participants achieved an accuracy of 74% in their personal choice (Fig. 4.2a), and participants reporting higher confidence in their personal choice were also more accurate (Fig. 4.2b; $\beta = 3.82$, $CI = [3.35, 4.28]$). Participants were thus—at least partly—aware of the quality of their personal information. We fitted a Two-Stage Dynamic Signal Detection model (2DSD; Pleskac and Busemeyer, 2010) to the choice, RT, and confidence data from the personal phase (see Supplementary Results and Discussion). The close correspondence between the model and the data suggests that a drift-diffusion process is a good description of the decision process during this stage of the experiment.

With an average accuracy of 79%, participants' choices during the social phase, where they had the opportunity to wait for social information before choosing again, were more accurate (Fig.




Figure 4.2. Choice accuracy and the relationship between personal accuracy and confidence. (a) The accuracy of the personal and social choices. Individuals, on average, achieved a higher decision accuracy during the social choice as compared to the personal choice. Each line connects a participant's average accuracy during the personal and social choice ($n=141$ participants). (b) Participants reporting a higher confidence in their personal choice were more likely to be correct in their personal choice. The points and error bars reflect the mean and the 95% credible intervals of the posterior distribution from the Bayesian logistic regression model.

4.2a; $\beta = 0.3$, $CI = [0.20, 0.39]$). The reported level of confidence in their personal choice predicted their likelihood to improve (Fig. 4.3a; $\beta = -4.27$, $CI = [-4.88, -3.68]$): participants reporting the lowest confidence level improved in more than 15% of trials; whereas the most confident, in less than 1% of trials. Why do unconfident participants achieve such higher gains from the social process? There are two mechanisms underlying this. First, participants reporting lower confidence waited longer before making a decision during the social phase (Fig. 4.3b; $\beta = -4.86$, $CI = [-5.22, -4.5]$). Second, participants partly adopted the decisions of others (Fig. 4.3c; $\beta = 0.62$, $CI = [0.57, 0.67]$): the larger the majority for the opposing option, the more likely participants were to change their decision. Individuals rarely changed their minds if the majority agreed with their personal decision. As Supplementary Figure C1 shows, participants followed both correct and incorrect majorities, highlighting the importance of the accuracy of early-deciding participants for triggering positive/negative information cascades. Figure 4.3d shows the consequences of these patterns: participants whose personal choices were accurate (and confident) tended to respond early in the social phase, whereas those whose choices were inaccurate (and unconfident) tended to wait longer, as illustrated by the downward trend of the blue dots (slope: $\beta = -0.16$, $CI =$

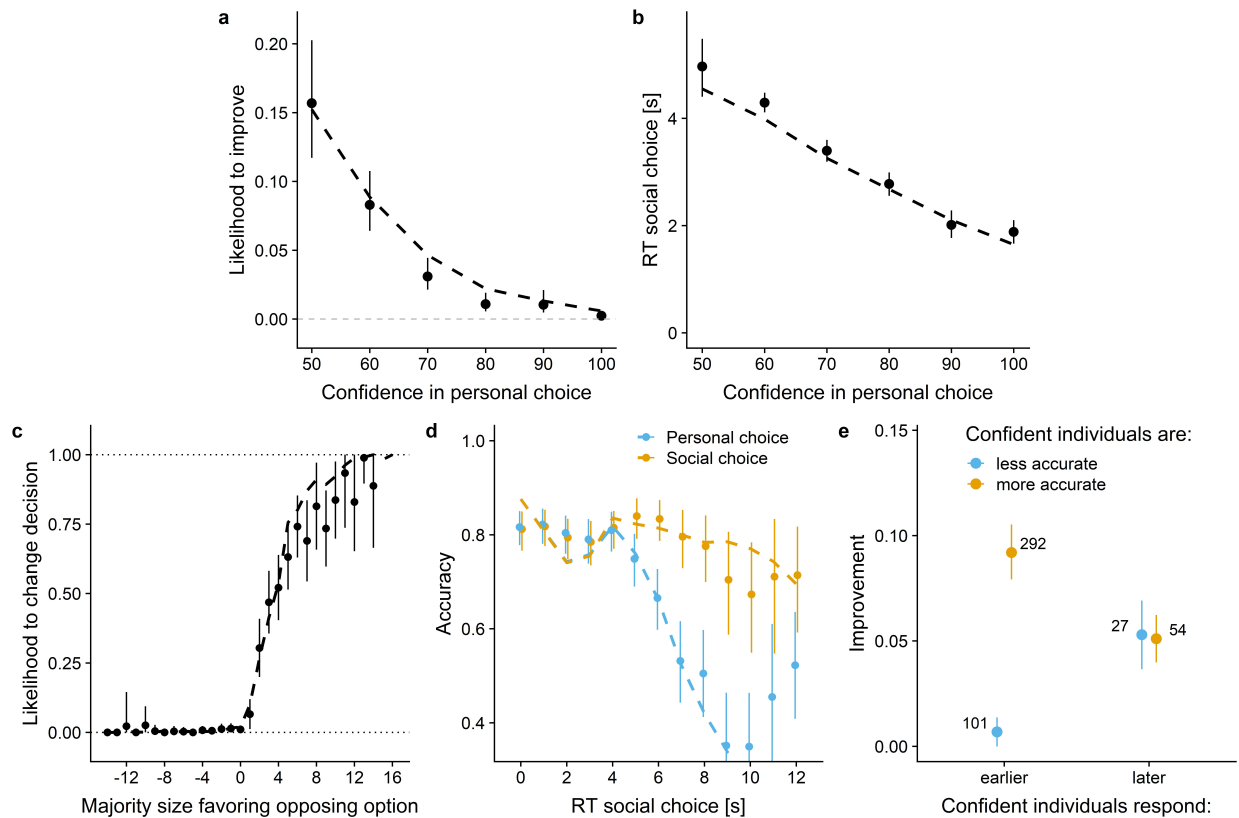


Figure 4.3. Empirical results and predictions of the social DDM. Participants reporting higher confidence in their personal choice (a) improved less and (b) responded earlier during the social choice. (c) The larger the majority favouring the opposing option, the more likely participants were to change their decision. (d) The choices of participants who responded later in the social choice were less accurate in the personal choice (declining blue dots) but improved more in the social choice (indicated by the increasing difference between blue and yellow dots at later RTs). For visualization purpose, RTs are binned by rounding to the closest integer. RTs greater than 13 seconds (less than 1%) were assigned to the 12 seconds bin. (a–d) The dashed lines show the choices and RTs predicted by the social DDM, accurately capturing all relationships. For frequency distributions, see Supplementary Figure C2. (e) Participants improved most when more confident individuals were more accurate (yellow dots) and responded earlier. Numbers indicate the number of trials. For all panels, the points and error bars depict the mean and the 95% credible intervals of the posterior distribution of the Bayesian regression model.

$[-0.18, -0.14]$). The latter participants increased their accuracy during the social phase through social influence, as illustrated by the higher yellow dots compared to the corresponding blue dots at higher RTs (interaction: $\beta = 0.11$, $CI = [0.09, 0.13]$).

Participants in groups thus self-organized according to information quality, with confident and accurate participants deciding early, thereby providing high-quality information for the less confident and less accurate participants, who decided later. This beneficial self-organization depended on two crucial aspects: (i) a positive relationship between confidence and accuracy of personal choice across group members, and (ii) a negative relationship between confidence and RT during

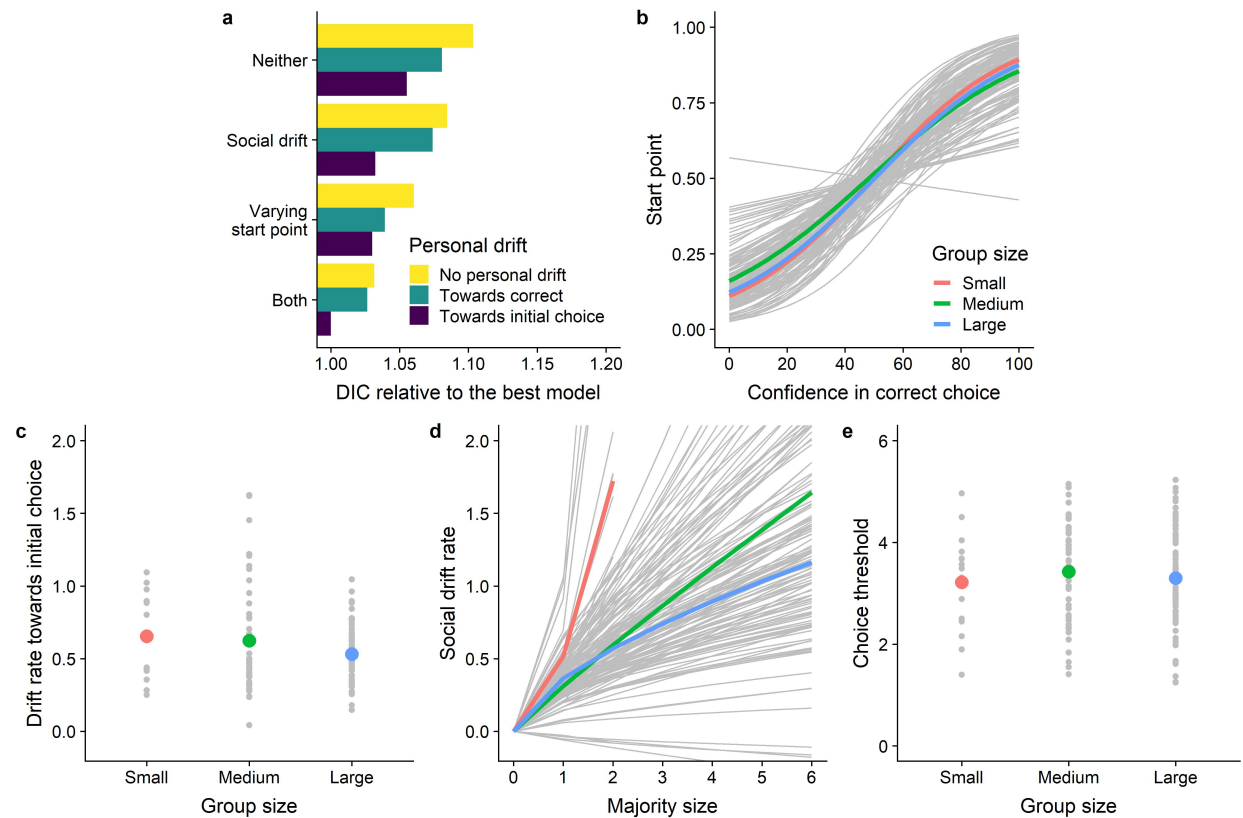


Figure 4.4. Model comparison and individual- and group-level fittings of the social DDM for different group sizes. (a) The deviance information criterion (DIC) values of all models relative to the model with the lowest DIC. The model with the lowest DIC (i.e., preferred model) features a (i) confidence-dependent start point, (ii) drift towards the initially chosen option, and (iii) social drift. (b) Participants reporting higher confidence in the correct/incorrect choice started closer to the correct/incorrect decision threshold at $y=1/0$. (c) Evidence tended to drift towards the choice threshold of the option chosen during the personal phase. (d) The larger the majority favouring an option, the more strongly participants drifted towards the choice threshold favoured by the majority. Participants in smaller groups had a stronger drift given the same majority size. (e) The choice threshold θ , reflecting a participant’s willingness to wait for social information, did not differ between group sizes. Grey lines/dots represent individual-level fittings; coloured lines/dots, the estimates on a group size-level.

the social choice phase. As Figure 4.3e illustrates, groups showed the highest improvement when both conditions were met, and this occurred in the majority of trials. Improvement was credibly lower for all other conditions (Supplementary Table C1).

Model results: the cognitive mechanisms driving self-organization

To understand the processes leading to the self-organization of groups, we need to understand the cognitive mechanisms underlying individuals’ dynamic integration of personal and social information over time. To this end, we developed the social DDM (Fig. 4.1a; Table 4.1), which allowed us to test competing hypotheses on how participants integrate personal and social information over

time. We examined three model features: (A) Individuals base their start point on their personal decision and reported confidence. (B) When participants start drifting, they drift towards the correct option, their initially chosen option, or neither of the two. (C) When social information becomes available, participants drift towards the option favoured by the majority. We tested several candidate models composed of various combinations of these three features, and used the deviance information criterion (DIC; Spiegelhalter et al. (2002)) to compare their performance. Figure 4.4a shows the models' DIC values relative to that of the best model (see also Supplementary Table C2). In the following, we present the results of the model with the lowest DIC (see Supplementary Table C3 for parameter estimates). Finally, to test how the cognitive mechanisms were affected by group size, we compared the different group sizes (Supplementary Table C4).

Individuals incorporate personal information via start point and drift rate

Participants incorporated their personal information (i.e., personal choice and confidence) during the social decision process in two distinct ways. First, consistent with current models of choice and confidence judgements (Moran et al., 2015; Pleskac and Busemeyer, 2010; Yu et al., 2015), they shifted their start point towards their initially chosen option: Individuals who reported higher confidence in the [in]correct option started closer to the threshold of the [in]correct option (Fig. 4.4b; small: $a = 4.20$, CI = [3.11, 5.35]; medium: $a = 3.42$, CI = [2.81, 4.07]; large: $a = 3.90$, CI = [3.46, 4.37]). This implies that individuals with high confidence in their personal choice were more likely to decide in favour of this option and to do so fast. Second, participants drifted towards the threshold of their initially chosen option (Fig. 4.4c; small: $\delta_p = 0.65$, CI = [0.45, 0.86]; medium: $\delta_p = 0.62$, CI = [0.50, 0.75]; large: $\delta_p = 0.53$, CI = [0.47, 0.59]). Both processes were independent of group size (Supplementary Table C4). To sum up, across all group sizes, highly confident participants started close to the choice threshold of their initially chosen option and, on top of that, drifted towards that option, whereas participants with low confidence started out unbiased (i.e., in the middle between the thresholds).

Individuals incorporate social information via drift rate

We found that the drift rates were credibly influenced by the majority (Fig. 4.4d). The larger the majority favouring an option, the more strongly participants drifted towards that option. The shape of the relationship between majority size and social drift rate (the q parameter) differed between group sizes (small vs. medium: $q = 0.82$, CI = [0.22, 1.44]; medium vs. large: $q =$

0.27, $CI = [0.14, 0.41]$). In small groups, the drift rate increased exponentially with increasing majority size. In larger groups, each additional individual voting for the majority had less impact than the preceding one, and this function followed a concave shape. Accordingly, the influence of a single individual was larger in small groups than in large groups. Comparing the strength of the personal drift rate (i.e., towards the choice threshold of the initially chosen option) to the social drift rate (i.e., towards the option favoured by the majority) showed that a majority of approximately two is required to counteract an individual's tendency to drift towards the choice threshold reflecting their initial choice. This highlights participants' tendency to give personal information more weight than social information. Corroborating this finding, Figure 4.3c shows that a majority of approximately four participants in favour of the opposing option is required to induce a 50% likelihood of changing a participant's decision. Finally, we found that participants' willingness to wait for social information, captured by the threshold parameter θ , did not differ between group sizes (Fig. 4.4e).

Model predictions: the social DDM captures the self-organizing dynamics

Importantly, the model described above was able to recover all the key features of the dynamics of the social decision making process. The dashed lines in Figure 4.3 show the model predictions of the social DDM. In line with the empirical data, the social DDM predicts that unconfident participants wait longer before making a decision (Fig. 4.3b), that individuals are increasingly likely to follow the majority as the size of that majority increases (Fig. 4.3c), and that participants whose personal choices were inaccurate wait longer and improve more during the social phase (Fig. 4.3d). As a result, participants with low confidence in their personal choice improved most (Fig. 4.3a). We investigated the validity of the model with a parameter recovery analysis (see Supplementary Information). For all parameters, the generating and recovered parameters were highly correlated, implying that each parameter describes a distinct mechanism. Further, all recovered parameter estimates were close to the generating parameters, affirming the validity of the magnitude of the parameter estimates as captured by the social DDM (Supplementary Fig. C3).

Discussion

We have shown that the behaviour of individuals in a social sequential decision making task can be described by an evidence accumulation process whereby personal and social information is integrated until a decision is made, formalized by the social DDM. The model accurately predicts

decision time and choice by taking personal information, social information, and the willingness to wait for social information into account. It successfully captured all the interrelationships of the key behavioural results of the social phase, thereby revealing the cognitive underpinnings of the group-level self-organization according to information quality. Measuring how individuals process personal and social information affords a deeper understanding of how individuals in a social environment cope with the complex problem of evaluating personal information, how they time their decision, and incorporate social information.

During the social decision making process, individuals incorporated personal information in two ways: at the start of the process, they adjusted their subjective level of evidence to their confidence (i.e., they adjusted their start point), and during the process, they reinforced their ‘belief’ in their original choice over time (i.e., they drifted towards the decision threshold of their personal choice). We also found evidence for such ‘belief reinforcement’ over time in the personal phase (see 2DSD model analysis in the Supplementary Information). The reinforcement of initial beliefs can potentially have a large influence in real-world social choices. Because individuals generally gather personal information before receiving social information, reinforcement of initial beliefs can lead to situations where even strong counterfactual social information may no longer prove persuasive (i.e., confirmation bias; Klayman, 1995; Koriat et al., 1980; Nickerson, 1998). Many studies have found that individuals indeed weight personal information more strongly than social information, a phenomenon called egocentric discounting (e.g., Jayles et al., 2017; Larrick and Soll, 2006; Tump et al., 2018; Yaniv and Kleinberger, 2000). In almost all of these studies, participants made a personal judgement before receiving social information. When the order was reversed, the influence of social information indeed increased (Koehler and Beaugard, 2006). Our finding of belief reinforcement provides a compelling explanation for egocentric discounting, simply by providing personal information first. Future studies could test whether increasing the length of the delay between personal choice and provision of social information reduces the influence of social information, as predicted by the social DDM.

When looking at how social information entered the evidence accumulation process, we found that individuals incorporated social information by drifting towards the decision threshold favoured by the majority. The larger the majority size, the more strongly individuals drifted towards that majority choice. For medium- and large-sized groups, the relationship followed a concave power function, where each additional individual voting for the majority choice had less additional impact on the drift rate. Such saturating influence is consistent with the findings of earlier studies

(Asch and Guetzkow, 1951; Bond, 2005; Latané, 1981; Milgram et al., 1969). In groups of three, the relationship followed an exponential function. Weighting single choices less with increasing group size is probably an adaptive strategy: In larger groups, waiting for further decisions avoids confirming fast, but wrong, choices, as others can still correct initial mistakes. In small groups, fast but wrong choices will also occur, but since there are few others to correct those choices, there is little point in delaying a response via a reduced social drift rate.

The social DDM can also characterize other features of the dynamics of the social decision making process. Beyond capturing how social information impacts the accumulation of evidence, it also captures an individual's willingness to wait for social information via the threshold parameter θ . Thus, the model is able to distinguish, for instance, between individuals who are sensitive to majorities but unwilling to wait for social information and individuals who may be interested in observing the decisions of others but put more weight on their own personal information. The capacity to unify these different facets of social decision making within a single theoretical framework is a long-standing goal of social decision making in the areas of collective animal behaviour (Deneubourg et al., 1990; Sumpter and Pratt, 2008) and social psychology (Latané, 1981). Future studies could investigate the interrelationships between the different parameters, and potential links to established personality measures.

Previous studies have provided evidence for both positive information cascades, such as knowledgeable individuals leading others to resources or safety (Dyer et al., 2008; Kurvers et al., 2015; Stroeymeyt et al., 2011; Watts et al., 2016), and for negative ones, such as the spread of fake news, mobbing, or stampedes (Giraldeau et al., 2002; Bikhchandani et al., 1998; Raafat et al., 2009). Here, we have shown the importance of two key aspects promoting positive information cascades. First, a positive confidence/accuracy relationship across group members. In many contexts, confidence is a valid cue for accuracy (Freund and Kasten, 2012; Hertwig, 2012; Bahrami et al., 2012). The strongest association of confidence and accuracy across group members arises when all individuals are more confident when they are more accurate and when their confidence scales are well aligned (i.e., a given level of confidence implies the same level of accuracy across individuals; see also Marshall et al., 2017; Bang and Frith, 2017). The second key aspect promoting positive information cascades is a negative relationship between confidence and RT, meaning that more confident individuals respond faster. Several mechanisms in the social DDM can influence this relationship—for example, how individuals adjust their start point depending on their confidence. If confident individuals do not start closer to a decision threshold, they are not expected to respond

earlier. Also, interindividual differences in model parameters such as choice thresholds or personal and social drift rate can negatively impact the confidence–RT relationship.

The quality of information cascades is shaped by the relationship between accuracy and response time, whereby it is crucial for positive information cascades that accurate individuals respond faster than inaccurate individuals. The social DDM framework allows us to predict the quality of information cascades on the basis of individual or task characteristics. For example, if individuals differ in their ability to solve a task (e.g., individual differences in drift rates), those with higher ability are expected to make faster, more accurate decisions than the less competent ones, triggering positive information cascades. In contrast, when individuals differ systematically in their speed–accuracy tradeoff (e.g., differences in threshold separation; Chittka et al., 2009; Ratcliff et al., 2016), and groups harbour both fast, but inaccurate individuals and slow, but accurate individuals, we expect relatively many fast errors, triggering negative information cascades.

Because the DDM has been successful in accounting for behavioural phenomena across a wide range of tasks, our extension to social environments opens up new possibilities for studying a range of social and collective phenomena. It makes it possible to measure how individuals combine personal and social information and time their decisions whenever decisions are made sequentially and the choices are—at least partially—observable by others. We hope future work will apply and extend the social DDM to areas such as dynamics in consumer preferences (Chen, 2008), emergency evacuations (Moussaïd et al., 2016), and social media (Vosoughi et al., 2018), or to areas of animal social and collective behaviour such as predator detection and mate choice (Danchin et al., 2004).

Methods

Experimental procedure

Participants were 141 students from Wageningen University (the Netherlands) and the University of Bielefeld (Germany). Participants were divided into 16 groups, with group size ranging from small (3 individuals; 5 groups) to medium (7–10 individuals; 6 groups), to large (15–17 individuals, 5 groups; see also Supplementary Table C5). Prior to participation, each participant signed an informed consent form. Each group of individuals was seated on chairs facing a large screen. They were confronted with the following binary decision task: individuals briefly (for 2 seconds) observed an image of a shoal of 72 stylized fish (tuna and sharks aligned in an 8 x 9 grid; see Fig. 4.1b). Participants were instructed to choose “Escape” if there were five or more sharks and

“Stay” if there were four or fewer. The number of sharks present was three, four, six, or seven, and each number was repeated ten times, resulting in 40 trials. Treatment order was randomized. After observing a shoal of fish, individuals had five seconds to report their personal decision and an additional five seconds to report their confidence in their personal decision. Participants were instructed to use confidence as the subjective probability of being correct on a scale from 50% to 100%. In the subsequent social phase, participants made a second decision on the same image. During this phase, they received social information in the form of the number of group members who had already decided on a particular option, displayed on the screen. The social information was first updated after three seconds and then iteratively every two seconds (i.e., at sec 3, 5, 7, 9, . . . 19). The social phase lasted 20 seconds. A countdown timer on the screen indicated the remaining choice time. Participants made all decisions using a wireless keypad. Afterwards, we provided feedback on the correct choice. Participants received 0 points for an incorrect decision and 100 points for a correct decision. To avoid a scenario in which all participants waited until the last second for social information, we introduced a small cost of one point per second for correct decisions during the social phase. The members of each group with the highest payoff got a small reward in kind. Prior to the 40 study trials, participants completed two test trials to familiarize themselves with the procedure. These results were excluded from the analyses.

Statistical analysis

We used Bayesian hierarchical generalized linear models with the “brms” package (Bürkner et al., 2017) to analyse the empirical data in R (R Core Team, 2019). The parameter estimates were generated by running five Markov Chain Monte Carlo (MCMC) simulations in parallel with 5,000 iterations, of which the first 2,500 were discarded as burn-in to reduce autocorrelations. To analyse the difference in the accuracy of personal and social choices (Fig. 4.2a), we fitted choice correct (yes/no) as a binary response variable and type of choice (personal/social) as a population-level effect (i.e., fixed effect). In this model (and all following models, unless stated otherwise), we included individual and group identity as group-level effects (i.e., mixed effects). We ran separate models to investigate how confidence related to (i) personal accuracy (Fig. 4.2b), (ii) likelihood to improve (Fig. 4.3a), and (iii) RT during social choice (Fig. 4.3b). ‘Personal accuracy’ (correct/incorrect) and ‘likelihood to improve’ (yes/no) were fitted as binomial response variables and ‘RT during social choice’ as an exponentially modified Gaussian (ex-Gaussian) distributed response variable. Confidence was included as a population-level effect in all three models. To investigate

whether the majority size affected the likelihood of an individual changing its decision (Fig. 4.3c), we fitted the likelihood to change the decision as a binary response variable (yes/no) and majority size favouring the opposing option as a population-level effect. To analyse the relationship between RT in the social phase and accuracy of personal and social choices (Fig. 4.3d), we used decision correct (yes/no) as a binary response variable and type of choice (personal/social) in interaction with RT as a population-level effect.

To investigate how the interrelationships between confidence, accuracy, and RT affected improvement (Fig. 4.3e), we first calculated—for each group and trial—the Spearman’s correlation coefficients of confidence and accuracy as well as of confidence and RT. We converted these coefficients into dichotomous variables, with the correlation coefficient being either 0 and above or below 0. We excluded trials in which all individuals reported identical choices or confidences, because it was impossible to calculate correlation coefficients for these. We treated all four possible combinations of correlations as different levels of a single factor. We included the factor as a population-level effect and improvement as response variable. In this model, group identity was the only group-level effect. As statistical summary, we report the mean of the posterior distributions and the 95% credible intervals (CI). See Supplementary Table C1 for the results of the regression models. To visualize the results (Fig. 4.2 and Fig. 4.3), while accounting for the hierarchical structure of the data, we re-ran the regression models, treating the continuous variables as categorical data. Unless stated otherwise, the points and error bars reflect the mean and the 95% CI of the posterior distribution. Visual inspection of the Markov chains and the Gelman Rubin statistic (\hat{R}) indicated that all Markov chains converged.

Social DDM: Model parameter estimation

To understand the dynamics of the social phase, we developed the social DDM (Fig. 4.1a, Table 4.1). The model features decisions with variable drift rates, in order to obtain choice and RT predictions. We calculated the probability density function of RTs and associated choice probabilities of the drift-diffusion process by implementing an extended version of a Markov chain approach (Diederich, 1997; Diederich and Busemeyer, 2003) in R (R Core Team, 2019). A detailed description of how to implement the Markov chain approach can be found in Diederich and Busemeyer (2003).

The model assumes that the state space of the decision maker’s evidence L is ranging from the lower decision threshold $-\theta$ (reflecting the wrong decision) to the upper threshold θ (reflecting

the correct decision) with a step size of Δ and k being the number of steps to reach the decision threshold from a neutral start point:

$$L = [-k\Delta, -(k-1)\Delta, \dots, -\Delta, 0, \Delta, \dots, (k-1)\Delta, k\Delta]; \quad (4.4)$$

where $\theta = k\Delta$.

Each time step h the evidence states change with probabilities given by a $m \times m$ transition probability matrix P , with $m = \frac{2 \times \theta}{\Delta} + 1$. The elements $p_{1,1} = 1$ and $p_{m,m} = 1$ are the two absorbing states and reflect the decision thresholds. The other elements of P with $1 < i < m$ are:

$$p_{i,j} = \begin{cases} \frac{1}{2\alpha} \left(1 - \frac{u}{\sigma^2} \sqrt{h}\right) & \text{if } j=i-1 \\ \frac{1}{2\alpha} \left(1 + \frac{u}{\sigma^2} \sqrt{h}\right) & \text{if } j=i+1 \\ 1 - \frac{1}{\alpha} & \text{if } j=i \\ 0 & \text{otherwise} \end{cases} \quad (4.5)$$

with σ^2 being the diffusion coefficient and $u = \delta_p + \delta_s$ the total drift rate, whereby δ_p and δ_s are drift rates reflecting the accumulation of personal and social information, respectively (see Table 4.1). The parameter $\alpha > 1$ improves the approximation of the continuous time process. We set $\alpha = 1.3$, $\sigma^2 = 1$ and $h = 0.005$. The transition probability matrix in its canonical form:

$$P = \left[\begin{array}{c|c} P_I & 0 \\ \hline R & Q \end{array} \right] = \begin{pmatrix} & 1 & m & & 2 & 3 & \dots & m-2 & m-1 \\ 1 & 1 & 0 & & 0 & 0 & \dots & 0 & 0 \\ m & 0 & 1 & & 0 & 0 & \dots & 0 & 0 \\ \hline 2 & p_{12} & 0 & & p_{22} & p_{23} & \dots & 0 & 0 \\ 3 & 0 & 0 & & p_{32} & p_{33} & \dots & 0 & 0 \\ 4 & 0 & 0 & & 0 & p_{43} & \dots & 0 & 0 \\ \vdots & \vdots & \vdots & & \vdots & \vdots & \ddots & \vdots & \vdots \\ m-3 & 0 & 0 & & 0 & 0 & \dots & p_{m-3,m-2} & 0 \\ m-2 & 0 & 0 & & 0 & 0 & \dots & p_{m-2,m-2} & p_{m-3,m-1} \\ m-1 & 0 & p_{m-1,m} & & 0 & 0 & \dots & p_{m-1,m-2} & p_{m-1,m-1} \end{pmatrix} \quad (4.6)$$

With P_I being a 2×2 matrix with the two absorbing states and R a $(m-2) \times 2$ matrix containing the transition probabilities that eventually lead to the absorbing states in a single transition. Q is a

$(m-2) \times (m-2)$ matrix including the remaining transition probabilities. The initial evidence state of the process is represented by Z an $m-2$ vector containing the initial probability distribution. The initial start point β is a function of confidence and choice and relative to the upper (correct) threshold:

$$\beta = \frac{1}{1 + e^{-a(C-b)}}; \quad (4.7)$$

with a and b being free parameters. C is the reported confidence in the correct choice and is scaled from zero (i.e., highly confident and wrong) to one (i.e., highly confident and correct). We set the distribution of the initial evidence states by $Z_{\beta^*} = 1$, with $\beta^* = \beta(m-3) + 1$. Because β^* is not always an integer, we avoid rounding errors by giving most probability mass $1 - (\beta^* - \text{round}(\beta^*))$ to $Z_{\text{round}(\beta^*)}$ and the rest $\beta^* - \text{round}(\beta^*)$ to the closest integer of β^* . For example, if the process starts unbiased (i.e., $\beta = 0.5$) and $m = 7$ then $\beta^* = 3$ and $Z = [0, 0, 1, 0, 0]$. However, if $\beta = 0.55$, then $\beta^* = 3.2$ and $\beta^* - \text{round}(\beta^*) = 0.2$ and therefore $Z = [0, 0, 0.8, 0.2, 0]$. We account for variable drift rates by updating the transition probabilities of Q at $t = (3, 5, 7, 9, \dots, 19)$ seconds, reflecting the iterative updated social information (see Experimental procedure). With Q_n containing the transition probabilities at time point $t = nh$, we can calculate the probability of choosing the correct or wrong option after n time steps:

$$[Pr(\text{correct}|n), P(\text{wrong}|n)] = Z \times Q_1 \times Q_2 \times Q_3 \dots Q_n \times R - \tau \times t_{min}, \quad (4.8)$$

with τ being the non-decision time relative to the fastest response of the individual t_{min} . By varying the transition probabilities of Q_n with changing δ_s we are able to account for varying social information over time.

Integrating the social DDM into a Bayesian estimation technique, namely a Differential-Evolution-MCMC algorithm, enables us to sample posterior probability densities of the model parameters (see Table 4.1). The Differential-Evolution-MCMC is an extension of the Metropolis-Hastings algorithm where proposals are generated by taking the Markov states of parallel computed chains into account (Ter Braak, 2006; Turner et al., 2013). To estimate the effect of group size while controlling for individual differences, we used a hierarchical framework. Each parameter was fitted on an individual level but was simultaneously informed by a higher order group-level prior, a normal distribution described by two hyper parameters (i.e., mean and variance), which were informed by the individual fittings. To estimate the posterior probability densities we ran 24 chains in parallel, each with a chain length of 20,000 including a burn-in period of 10,000 and a thinning factor of 10 to reduce autocorrelations. The tuning parameter (γ) was set to $= 2.38/\sqrt{2d}$,

with d being the dimensionality of the posterior, which was $d = 2$ for the hyper parameters and $d = 7$ for the individual parameters (see Ter Braak, 2006; Turner et al., 2013). To further improve the mixing of the parallel chains, we included deterministic and probabilistic (i.e., relying on the Metropolis–Hastings probability) migration steps where chain-states are swapped across parallel chains (Turner et al., 2013). We performed the deterministic migration step with a probability of 5% where we first determine a random number of $n = 2, 3, \dots, 24$ chains and then sample n chains without replacement. We then swap the parameter set in a cyclic fashion where the set of the first sampled chain moves to the second, the second to the third and so on, until the last set moves to the first set. A deterministic migration step strongly improves the mixing behaviour of chains but does not resolve the frequent problem of Differential-Evolution-MCMC algorithms that outlier chains hardly converge. We, therefore, additionally implemented a probabilistic migration step which was carried out with a probability of 10%. For the probabilistic version we swapped proposal states instead of accepted states between chains which therefore still relied on the Metropolis–Hastings probability to be accepted. Thereby, we sampled two parallel chains and interchanged a single random parameter state.

We used the social DDM to compare competing hypotheses on how individuals integrate personal and social information. More specifically, we examined three model features: (A) Individuals base their start point on their personal choice and reported confidence. (B) Individuals drift towards the correct option, their initially chosen option, or neither of the two. (C) Individuals drift towards the option favoured by the majority. We compared the performance of models composed of the various combinations of these three features using the deviance information criterion (DIC; Spiegelhalter et al., 2002). To investigate the effect of group size on the collective dynamics, we categorized groups as small (3 individuals), medium (7–10), or large (15–17), and fitted the parameters separately for each group size. As a statistical summary, we report the mean of the posterior distributions and the 95% CI. We excluded all observations for which personal choice, social choice, confidence, or RT of social choice were missing (~8%) and if the RT was below 0.1 second (~6%).

Social DDM: predictions

To analyse the predictions (i.e., choices and RTs) of the social DDM, we generated decisions by sampling from the probability density functions produced by the model using the mean of the individual-level posterior distribution as model estimates. The probability density function was

computed for each individual and trial by taking into account the individual model estimates, the personal choice, the reported confidence, and the social information observed by the individual at a given trial. We then sampled 10 choices and RTs to account for stochasticity. The model predictions are shown as dashed lines in Figure 4.3.

References

- Anderson, L. R. and Holt, C. A. (1997). Information cascades in the laboratory. *The American Economic Review*, 87(5):847–862.
- Asch, S. E. and Guetzkow, H. (1951). *Effects of group pressure upon the modification and distortion of judgments*. Carnegie Press, Oxford, England.
- Bahrami, B., Olsen, K., Bang, D., Roepstorff, A., Rees, G., and Frith, C. (2012). What failure in collective decision-making tells us about metacognition. *Philosophical Transactions of the Royal Society B*, 367(1594):1350–1365.
- Banerjee, A. V. (1992). A simple model of herd behavior. *The Quarterly Journal of Economics*, 107(3):797–817.
- Bang, D. and Frith, C. D. (2017). Making better decisions in groups. *Royal Society Open Science*, 4(8):170193.
- Battaglini, M. (2005). Sequential voting with abstention. *Games and Economic Behavior*, 51(2):445–463.
- Bikhchandani, S., Hirshleifer, D., and Welch, I. (1998). Learning from the Behavior of Others: Conformity, Fads, and Informational Cascades. *Journal of Economic Perspectives*, 12(3):151–170.
- Bond, R. (2005). Group size and conformity. *Group Processes & Intergroup Relations*, 8(4):331–354.
- Brown, S. D. and Heathcote, A. (2008). The simplest complete model of choice response time: Linear ballistic accumulation. *Cognitive Psychology*, 57(3):153–178.
- Bürkner, P.-C. et al. (2017). brms: An R package for Bayesian multilevel models using Stan. *Journal of Statistical Software*, 80(1):1–28.
- Busemeyer, J. R. and Diederich, A. (2002). Survey of decision field theory. *Mathematical Social Sciences*, 43(3):345–370.
- Busemeyer, J. R. and Townsend, J. T. (1993). Decision field theory: A dynamic-cognitive approach to decision making in an uncertain environment. *Psychological Review*, 100(3):432–459.
- Chamley, C. and Gale, D. (1994). Information revelation and strategic delay in a model of investment. *Econometrica*, 62(5):1065.
- Chen, Y.-F. (2008). Herd behavior in purchasing books online. *Computers in Human Behavior*, 24(5):1977–1992.
- Chittka, L., Skorupski, P., and Raine, N. E. (2009). Speed–accuracy tradeoffs in animal decision making. *Trends in Ecology & Evolution*, 24(7):400–407.
- Dall, S. R., Giraldeau, L.-A., Olsson, O., McNamara, J. M., and Stephens, D. W. (2005). Information and its use by animals in evolutionary ecology. *Trends in Ecology & Evolution*, 20(4):187–193.
- Danchin, E., Giraldeau, L.-A., Valone, T. J., and Wagner, R. H. (2004). Public information: from nosy neighbors to cultural evolution. *Science*, 305(5683):487–491.
- Deneubourg, J.-L., Aron, S., Goss, S., and Pasteels, J. M. (1990). The self-organizing exploratory pattern of the Argentine ant. *Journal of Insect Behavior*, 3(2):159–168.
- Diederich, A. (1997). Dynamic stochastic models for decision making under time constraints. *Journal of Mathematical Psychology*, 41(3):260–274.
- Diederich, A. and Busemeyer, J. R. (2003). Simple matrix methods for analyzing diffusion models of choice probability, choice response time, and simple response time. *Journal of Mathematical Psychology*, 47(3):304–322.
- Dyer, J. R., Johansson, A., Helbing, D., Couzin, I. D., and Krause, J. (2008). Leadership, consensus decision making and collective behaviour in humans. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 364(1518):781–789.

- Edwards, W. (1965). Optimal strategies for seeking information: Models for statistics, choice reaction times, and human information processing. *Journal of Mathematical Psychology*, 2(2):312–329.
- Faria, J. J., Krause, S., and Krause, J. (2010). Collective behavior in road crossing pedestrians: The role of social information. *Behavioral Ecology*, 21(6):1236–1242.
- Freund, P. A. and Kasten, N. (2012). How smart do you think you are? A meta-analysis on the validity of self-estimates of cognitive ability. *Psychological Bulletin*, 138(2):296–321.
- Galef, B. G. and Laland, K. N. (2005). Social learning in animals: empirical studies and theoretical models. *Bioscience*, 55(6):489–499.
- Gallup, A. C., Hale, J. J., Sumpter, D. J., Garnier, S., Kacelnik, A., Krebs, J. R., and Couzin, I. D. (2012). Visual attention and the acquisition of information in human crowds. *Proceedings of the National Academy of Sciences*, 109(19):7245–7250.
- Germar, M., Schlemmer, A., Krug, K., Voss, A., and Mojzisch, A. (2014). Social influence and perceptual decision making: A diffusion model analysis. *Personality and Social Psychology Bulletin*, 40(2):217–231.
- Giraldeau, L.-A., Valone, T. J., and Templeton, J. J. (2002). Potential disadvantages of using socially acquired information. *Philosophical Transactions of the Royal Society of London. Series B: Biological Sciences*, 357(1427):1559–1566.
- Gul, F. and Lundholm, R. (1995). Endogenous timing and the clustering of agents’ decisions. *Journal of Political Economy*, 103(5):1039–1066.
- Hertwig, R. (2012). Tapping into the wisdom of the crowd—with confidence. *Science*, 336(6079):303–304.
- Hertwig, R., Pleskac, T. J., Pachur, T., and The Center for Adaptive Rationality (2019). *Taming uncertainty*. MIT Press, Cambridge, MA.
- Jayles, B., Kim, H.-r., Escobedo, R., Cezera, S., Blanchet, A., Kameda, T., Sire, C., and Theraulaz, G. (2017). How social information can improve estimation accuracy in human groups. *Proceedings of the National Academy of Sciences*, 114(47):12620–12625.
- Klayman, J. (1995). Varieties of confirmation bias. *Psychology of Learning and Motivation*, 32:385–418.
- Koehler, D. J. and Bearegard, T. A. (2006). Illusion of confirmation from exposure to another’s hypothesis. *Journal of Behavioral Decision Making*, 19(1):61–78.
- Konovalov, A. and Krajbich, I. (2019). Revealed indifference: Using response times to infer preferences. *Judgment and Decision Making*, 14(4):381–394.
- Koriat, A., Lichtenstein, S., and Fischhoff, B. (1980). Reasons for confidence. *Journal of Experimental Psychology: Human Learning and Memory*, 6(2):107–118.
- Krajbich, I. and Rangel, A. (2011). Multialternative drift-diffusion model predicts the relationship between visual fixations and choice in value-based decisions. *Proceedings of the National Academy of Sciences*, 108(33):13852–13857.
- Kurvers, R. H., Wolf, M., Naguib, M., and Krause, J. (2015). Self-organized flexible leadership promotes collective intelligence in human groups. *Royal Society Open Science*, 2(12):150222.
- Laming, D. R. J. (1968). *Information theory of choice-reaction times*. Academic Press, Oxford, England.
- Larrick, R. P. and Soll, J. B. (2006). Intuitions about combining opinions: Misappreciation of the averaging principle. *Management Science*, 52(1):111–127.
- Latané, B. (1981). The psychology of social impact. *American psychologist*, 36(4):343–356.
- Leite, F. P. and Ratcliff, R. (2011). What cognitive processes drive response biases? A diffusion model analysis. *Judgment & Decision Making*, 6(7):651–687.
- Link, S. W. and Heath, R. A. (1975). Sequential theory of psychological discrimination. *Psychometrika*, 40(1):77–105.
- Mann, R. P. (2018). Collective decision making by rational individuals. *Proceedings of the National Academy of Sciences*, 115(44):E10387–E10396.
- Marshall, J. A., Brown, G., and Radford, A. N. (2017). Individual confidence-weighting and group decision-making. *Trends in Ecology & Evolution*, 32(9):636–645.
- Milgram, S., Bickman, L., and Berkowitz, L. (1969). Note on the drawing power of crowds of different size. *Journal of Personality and Social Psychology*, 13(2):79–82.
- Moran, R., Teodorescu, A. R., and Usher, M. (2015). Post choice information integration as a causal determinant of confidence: Novel data and a computational account. *Cognitive Psychology*, 78:99–147.
- Moussaïd, M., Kapadia, M., Thrash, T., Sumner, R. W., Gross, M., Helbing, D., and Hölscher, C. (2016). Crowd behaviour during high-stress evacuations in an immersive virtual environment. *Journal of The*

- Royal Society Interface*, 13(122):20160414.
- Nickerson, R. S. (1998). Confirmation bias: A ubiquitous phenomenon in many guises. *Review of General Psychology*, 2(2):175–220.
- Nosofsky, R. M. and Palmeri, T. J. (1997). An exemplar-based random walk model of speeded classification. *Psychological Review*, 104(2):266–300.
- Pfeffer, K. and Hunter, E. (2013). The effects of peer influence on adolescent pedestrian road-crossing decisions. *Traffic Injury Prevention*, 14(4):434–440.
- Pleskac, T. J. and Busemeyer, J. R. (2010). Two-stage dynamic signal detection: A theory of choice, decision time, and confidence. *Psychological Review*, 117(3):864–901.
- Pleskac, T. J., Cesario, J., and Johnson, D. J. (2018). How race affects evidence accumulation during the decision to shoot. *Psychonomic Bulletin & Review*, 25(4):1301–1330.
- R Core Team (2019). *R: A language and environment for statistical computing*. R Foundation for Statistical Computing, Vienna, Austria.
- Raafat, R. M., Chater, N., and Frith, C. (2009). Herding in humans. *Trends in Cognitive Sciences*, 13(10):420–428.
- Ratcliff, R. (1978). A theory of memory retrieval. *Psychological Review*, 85(2):59–108.
- Ratcliff, R. and McKoon, G. (2008). The diffusion decision model: theory and data for two-choice decision tasks. *Neural Computation*, 20(4):873–922.
- Ratcliff, R. and Smith, P. L. (2004). A comparison of sequential sampling models for two-choice reaction time. *Psychological Review*, 111(2):333–367.
- Ratcliff, R., Smith, P. L., Brown, S. D., and McKoon, G. (2016). Diffusion decision model: Current issues and history. *Trends in Cognitive Sciences*, 20(4):260–281.
- Shiller, R. J. (2002). Bubbles, human judgment, and expert opinion. *Financial Analysts Journal*, 58(3):18–26.
- Spiegelhalter, D. J., Best, N. G., Carlin, B. P., and Van Der Linde, A. (2002). Bayesian measures of model complexity and fit. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 64(4):583–639.
- Stone, M. (1960). Models for choice-reaction time. *Psychometrika*, 25(3):251–260.
- Stroeymeyt, N., Franks, N. R., and Giurfa, M. (2011). Knowledgeable individuals lead collective decisions in ants. *Journal of Experimental Biology*, 214(18):3046–3054.
- Sumpter, D. J. and Pratt, S. C. (2008). Quorum responses and consensus decision making. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 364(1518):743–753.
- Ter Braak, C. J. (2006). A Markov Chain Monte Carlo version of the genetic algorithm Differential Evolution: easy Bayesian computing for real parameter spaces. *Statistics and Computing*, 16(3):239–249.
- Toelch, U., Panizza, F., and Heekeren, H. R. (2018). Norm compliance affects perceptual decisions through modulation of a starting point bias. *Royal Society Open Science*, 5(3):171268.
- Tump, A., Wolf, M., Krause, J., and Kurvers, R. H. (2018). Individuals fail to reap the collective benefits of diversity because of over-reliance on personal information. *Journal of the Royal Society Interface*, 15(142):20180155.
- Turner, B. M., Sederberg, P. B., Brown, S. D., and Steyvers, M. (2013). A method for efficiently sampling from distributions with correlated dimensions. *Psychological Methods*, 18(3):368.
- Tversky, A. and Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biases. *Science*, 185(4157):1124–1131.
- Usher, M. and McClelland, J. L. (2001). The time course of perceptual choice: The leaky, competing accumulator model. *Psychological Review*, 108(3):550–592.
- Vosoughi, S., Roy, D., and Aral, S. (2018). The spread of true and false news online. *Science*, 359(6380):1146–1151.
- Voss, A., Rothermund, K., and Voss, J. (2004). Interpreting the parameters of the diffusion model: An empirical validation. *Memory & Cognition*, 32(7):1206–1220.
- Watts, I., Nagy, M., Burt de Perera, T., and Biro, D. (2016). Misinformed leaders lose influence over pigeon flocks. *Biology Letters*, 12(9):20160544.
- Welch, I. (2000). Herding among security analysts. *Journal of Financial Economics*, 58(3):369–396.
- Xiong, F. and Liu, Y. (2014). Opinion formation on social media: An empirical approach. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 24(1):013130.

- Yaniv, I. and Kleinberger, E. (2000). Advice taking in decision making: Egocentric discounting and reputation formation. *Organizational Behavior and Human Decision Processes*, 83(2):260 – 281.
- Yu, S., Pleskac, T. J., and Zeigenfuse, M. D. (2015). Dynamics of postdecisional processing of confidence. *Journal of Experimental Psychology: General*, 144(2):489–510.
- Zhang, J. (1997). Strategic delay and the onset of investment cascades. *The RAND Journal of Economics*, 28(1):188–205.
- Ziegelmeyer, A., My, K. B., Vergnaud, J.-C., Willinger, M., et al. (2005). Strategic delay and rational imitation in the laboratory. Unpublished, Max-Planck-Inst. for Research into Economic Systems.

Chapter 5

Adaptive decision rules in groups under asymmetrical error costs

Tump, A.N., Wolf, M., & Kurvers, R.H.J.M.

In Preparation

Abstract

Across a wide range of contexts, decision makers face the challenge of compromising between two kinds of errors: false positives and false negatives. If the costs of these errors are asymmetrical, individuals acting alone are known to develop a response bias to avoid the more costly error. However, how individuals in groups cope with asymmetrical costs is not well understood. We used a drift–diffusion model to study the decision-making process of groups facing asymmetrical error costs. In the model, individuals first gather personal information alone; in a second phase, they can aggregate additional social information to make a decision. Individuals can either decide early on, potentially influencing others, or wait for more social information. We combined this with a genetic algorithm approach to derive the optimal behavior. Our results confirm that, under asymmetrical costs, small cooperative groups (where individuals aim to maximize group payoff) evolve a response bias to avoid the costly error. Large cooperative groups, however, do not evolve a response bias, since the danger of response biases triggering false information cascades increases with group size. In small and large competitive groups (where individuals aim to maximize individual payoff), individuals evolve a higher response bias and wait longer for social information, thereby undermining the overall group performance. Our results have broad implications for understanding social dynamics in situations where error costs are asymmetrical, such as crowd panics and predator detection.

Acknowledgments: We thank Deborah Ain for editing the manuscript. R.H.J.M.K. acknowledges funding from the German Research Foundation (grant number: KU 3369/1-1).

Code Availability: The code to implement all analysis can be accessed in

https://github.com/alantump/socialDDM_Evolutionary_Algorithm

Author Contributions: Conceptualization: A.N.T., R.H.J.M.K., & M.W.; A.N.T. carried out the simulations; A.N.T. and R.H.J.M.K. drafted the manuscript with substantial input from M.W..

Introduction

Whether detecting predators, making medical diagnoses, or assessing security, decision makers face a crucial challenge. They have to categorise the world into one of two possible states: signal (e.g., a predator, disease, or threat) or no signal (Beauchamp and Ruxton, 2007; Green and Swets, 1966; Macmillan and Creelman, 2005; Swets et al., 2000). The true state of the world, however, is unknown, and must be inferred from noisy cues. This leads to a fundamental trade-off in decision theory, namely, compromising between increasing true positives, or “hits” (e.g., escaping a predator, referring a diseased person for treatment) and decreasing false positives, or “false alarms” (e.g., escaping in the absence of a predator, referring a healthy patient for treatment). A decision maker who is sensitive to cues that potentially indicate the signal increases hits, but at the expense of an increase in false alarms. A decision maker who is unresponsive to cues reduces false alarms, but also hits. In most of these binary decision-making contexts, the costs (or base rates) an individual faces are not symmetrical. Failing to identify a signal, such as a predator or disease (“miss”), is generally more costly than a false alarm (Johnson et al., 2013). It is therefore critical to take asymmetries into account in adaptive decision making (Johnson et al., 2013; Marshall et al., 2019; Mulder et al., 2012).

A substantial body of literature has investigated how individuals adjust their decision-making strategies to differences in cost (or base rate) asymmetries, showing that individuals are likely to develop a response bias that allow them to avoid the more costly error when faced with asymmetrical costs (Green and Swets, 1966; Maddox, 2002; Mulder et al., 2012; Ratcliff and McKoon, 2008; Swets et al., 2000). However, the way groups of decision makers should and actually do deal with asymmetrical costs has received far less attention. Given the widespread manifestations of collective decisions across biological systems—from bacteria to cells to groups of animals—and the widespread presence of asymmetrical error costs, this is an important knowledge gap, not least because individual response biases can have substantial bearings on collective responses. Wolf et al. (2013) studied how to optimally pool information in a collective decision-making scenario in which all individuals simultaneously indicated their decision. They showed that optimal decisions arise when individuals use a quorum threshold between the true and false positive rate of their group members. Building on this, Marshall et al. (2019) investigated how to optimally pool independent decisions for various error costs and base rates, showing that the use of quorums is extremely powerful for optimizing decision making across a broad range of environmental conditions (see

also Ben-Yashar and Nitzan, 1997).

These studies assume that all group members decide simultaneously, and that individuals have information about the group's average opinion. Both assumptions are unrealistic for decision-making processes in most biological systems, where individuals decide sequentially, with more knowledgeable (or confident) individuals making faster decisions (Kurvers et al., 2015; Stroeymeyt et al., 2011; Tump et al., 2019; Watts et al., 2016). These early decisions can influence later-deciding individuals, and, in extreme cases, early-deciding individuals can even trigger information cascades in which all individuals imitate early decisions (Anderson and Holt, 1997; Banerjee, 1992; Bikhchandani et al., 1998; Gallup et al., 2012). Understanding the dynamics of these cascades, in particular when and why they go wrong (e.g., crowd panics), is relevant for a wide range of social systems.

To study optimal decision making under asymmetrical costs, we developed a dynamic agent-based model that can account for a realistic decision-making scenario of sequential decision making. We modeled the decision process using a social drift–diffusion model (social DDM) and retrieved the adaptive—and therefore optimal—behavioral parameters using evolutionary algorithms. In the model, each decision maker first accumulates its own personal information about the state of the world, then enters a social phase. During the social phase, decision makers aggregate their personal information with social information (i.e., the choices of others). When a decision maker has gathered sufficient evidence (i.e., the decision threshold is exceeded), the decision is made. Because individuals both emit and receive social information, the system is highly dynamic and the final outcome is influenced by early responders and the initial evidence distributions among group members. We systematically varied the asymmetry in costs to study how this affects the evolution of several key decision-making parameters, such as the evolution of a response bias and the use of social information, across group size. We started with groups of individuals whose interests were completely aligned and selection operates on the mean group payoff (hereafter: “cooperative groups”). We then investigated how introducing competition between individuals shaped optimal behavior across different cost asymmetries and group sizes. In short, we found that individuals facing high asymmetrical costs in small cooperative groups optimized payoff by evolving a high bias towards the signal response. Large cooperative groups, however, do not evolve a bias, even when facing strong asymmetrical costs. Across all group sizes, adding competition increased the optimal bias while simultaneously reducing performance, showing how competition can impair collective benefits.

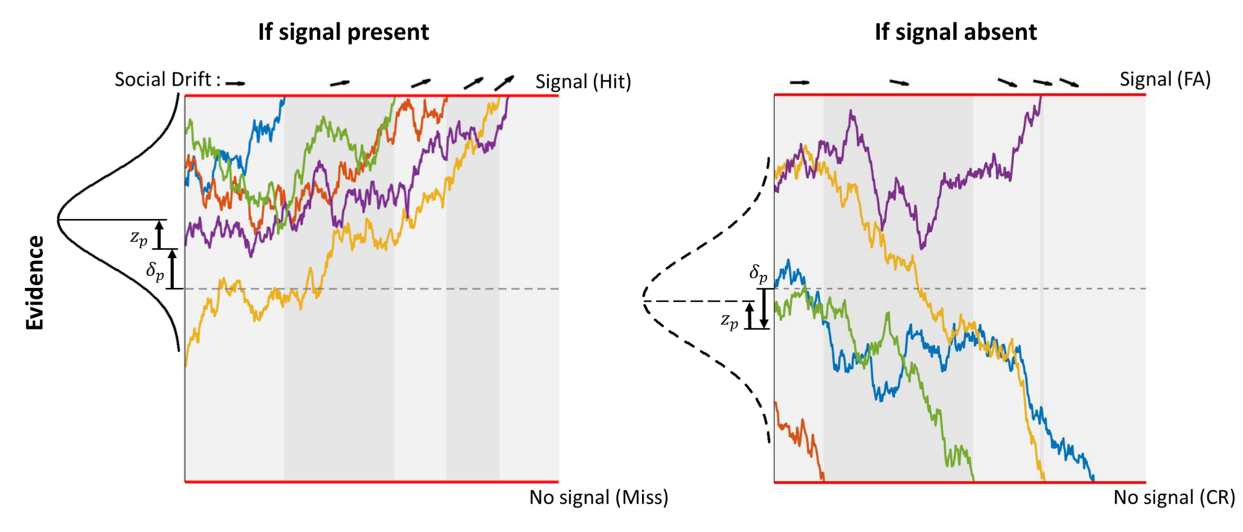


Figure 5.1. Illustration of the social DDM. Each decision maker, represented by a jagged line, has to decide whether a signal is present (left panel) or absent (right panel). If the signal is present, the individual can decide correctly (hit) or wrongly (miss). If the signal is absent, the individual can be correct (correct rejection, CR) or wrong (false alarm, FA). The start point of each individual depends on the information it gathered prior to the social process (δ_p) and on its bias (z_p). Here the bias is towards the decision boundary of the signal, implying that an individual is more likely to make a correct decision when the signal is present, but more likely to err when the signal is absent, reflecting the trade-off between increasing hits at the expense of increasing false alarms. At the start of the drift, no individual reached either decision boundary, implying that social information was absent. As individuals diffuse they hit a decision boundary and make a decision. Undecided individuals, in turn, start drifting towards the choice of the individuals that already decided, reflecting the process of social information use.

Methods

The social DDM. In the social DDM a group of individuals faces a signal detection task, choosing between two states of the world: signal present or signal absent. There are thus four possible decision outcomes: Individuals can correctly decide that a signal is present, correctly decide that a signal is absent, incorrectly decide that a signal is present, or incorrectly decide that a signal is absent. The decision-making process consists of two phases: first, a personal phase in which individuals independently sample information from the environment, then a social phase in which individuals decide on an option in the presence of others.

In the personal phase, each individual independently gathers information about the state of the world. The personal evidence accumulation process is formalized by a diffusion process (Gold and Shadlen, 2007; Pleskac and Busemeyer, 2010), whereby individuals start with an initial bias z_p and, over time, gather, on average correct evidence described by a positive drift rate δ_p . The total amount of evidence $L(t_p)$ at time point t_p is described by a normal distribution with a mean of

$$E[L(t_p)] = \begin{cases} \delta_p \times t_p + z_p, & \text{if signal is present} \\ \delta_p \times t_p - z_p, & \text{if signal is absent} \end{cases} \quad (5.1)$$

and a variance of

$$\text{Var}[L(t_p)] = \sigma_p^2 \times t_p; \quad (5.2)$$

with σ_p being the diffusion rate (σ_p and t_p are set to 1 for simplicity). The parameters δ_p and z_p change the state of evidence in distinct ways. A positive (negative) δ_p shifts the mean towards the decision boundary of the correct (wrong) decision; a positive (negative) z_p shifts the mean towards the decision boundary of the signal (no signal) option (see Fig. 5.1).

Next, individuals enter a social phase—also formalized as a drift–diffusion process. In this phase, individuals no longer sample independent information from the environment; instead, they can update their evidence based on the decisions of others. At each time step t the current state of evidence $L(t)$ is updated by sampling new evidence until a decision is made (i.e., the level of evidence reaches the decision boundary at $\pm \frac{\theta}{2}$):

$$L(t + t) = L(t) + \delta_s \times t + \sqrt{t} \times \epsilon, \quad (5.3)$$

where ϵ is white noise (i.e., the diffusion process) with a mean of zero and a variance of one. The social drift rate $\delta_s(t)$ describes the change in an individual’s drift rate depending on the decisions of others (i.e., the impact of social information) and is modeled proportional to the size of the majority $M(t)$ of individuals who already decided at time point t (e.g., see Bikhchandani et al., 1998; Tump et al., 2019):

$$M(t) = N^+(t) - N^-(t), \quad (5.4)$$

$$\delta_s(t) = s \times M(t), \quad (5.5)$$

where $N^+(t)$ and $N^-(t)$ are the number of individuals that decided the signal was present or absent, respectively, at time point t and s is scaling the strength of the social drift. This model implementation ensures that individuals with strong personal information will, on average, start closer to one of the decision boundaries and thus make faster decisions. This captures the observation that, in groups across many biological systems, individuals that are better-informed make faster decisions (Couzin et al., 2005; Kurvers et al., 2015; Reeb, 2000; Stroeymeyt et al., 2011). Because early-deciding individuals can influence the drift rate of undecided individuals, information

can flow from early-deciding, well-informed individuals to later-deciding, less-informed individuals, capturing the process of information cascades (Tump et al., 2019).

To study the effect of asymmetrical error costs, we implemented different payoff structures. Under symmetrical costs, an individual received one point for a correct decision (hit or correct rejection) and lost one point for a wrong decision (miss or false alarm). We modeled asymmetries in error costs by increasing the cost ratio of a miss compared to that of a false alarm (see The evolutionary algorithm), reflecting that missing a signal (e.g., a predator or disease) generally incurs a higher cost than does a false alarm. We also included a time cost of 0.05 points per second, which is subtracted from the payoff whether the choice is correct or not. This time cost reflects the ecologically valid benefit of making fast choices (Chittka et al., 2009).

The evolutionary algorithm. We embedded the social DDM into an evolutionary algorithm to derive the optimal behavior. Inspired by biological evolution, evolutionary algorithms allow fitness-maximizing parameter settings to evolve by exposing them to selection pressure and mutation (Hamblin, 2013). These algorithms are particularly suited for game-theory problems, where the optimal behavior of individuals depends on the behavior of others. Each individual had three genes, coding their parameter values for their bias z_p , boundary separation θ , and strength of the social drift s (Table 5.1). The parameters covered a wide range, ensuring the best solutions were included (range for bias: -0.5–2; boundary separation: 0.01–12; social drift: 0–2). To ensure that the end points of the simulations were independent of their starting conditions, we sampled the initial parameters from a beta distribution with the minimum and maximum scaled to the respective evaluated parameter range, whereby the mean of the beta distribution was sampled from a uniform distribution. In each generation, individuals were randomly and repeatedly (on average 10 times) sampled from the population; they performed the social DDM simulation as described above in the presence of other individuals (in different-sized groups; see below). Population size was fixed at 1,000 individuals across all conditions. After these simulations, individuals produced offspring based on their sum payoff, implemented via tournament selection: Three individuals were randomly sampled from the population and the individual with the highest payoff passed its genes to the next generation. This procedure was repeated 1,000 times. Finally, the genes of the new generation were exposed to mutation and crossover to ensure genetic variation. Crossover was implemented by swapping two genes between a focal and a randomly drawn individual with a probability of 0.05. For the mutation process, we added Gaussian noise to a gene with a mutation

Table 5.1. Description of model parameters. Underlined parameters evolve in the evolutionary algorithm.

Model feature	Parameter	Description
Start point	$L(t_p) \sim N(\delta_p \pm \underline{z_p}, \sigma_p)$	Parameters influencing the evidence gathered during the personal phase $L(t_p)$, which then served as the start point in the social phase. δ_p , $\underline{z_p}$ determine the mean and σ_p the variance of a normal distribution. A positive (negative) δ_p shifts the mean towards the correct (wrong) option, reflecting the amount of correct evidence gathered. A positive (negative) $\underline{z_p}$ shifts the mean towards the signal (no signal) response, reflecting a response bias.
Boundary separation	$\underline{\theta}$	The boundary separation determines how much evidence an individual accumulates before making a decision. Increasing the boundary separation increases the potential for social information use.
Social drift rate	$\delta_s(t) = \underline{s} \times M(t)$	\underline{s} determines how strong individuals incorporate social information by adjusting the strength of the social drift rate and the size of the majority of individuals who already decided for a particular option $M(t)$ at time point t .

probability of 0.02 and a standard deviation of 5% of the evaluated parameter ranges. These procedures were repeated for 1,000 generations, ensuring that populations converged to stable end points. We measured the evolved parameters by averaging the parameter values of the last 10 generations across eight populations.

We systematically varied three features to study their impact on the evolution of optimal parameter settings. First, we varied the group size (1, 5, 10, 20, and 50) in which individuals made decisions. Second, we varied the miss/false alarm cost ratio (1, 2, and 4). For this, we kept the average cost constant at 1 by increasing the cost of a miss from 1 to $1.\overline{33}$ to 1.6, and decreasing the cost of a false alarm from 1 to $0.\overline{66}$ to 0.4. Third, individuals received either a payoff based on their mean group payoff (cooperative scenario) or their own payoff (competitive scenario).

Performance evaluation. To gain a deeper understanding of the individuals' behavior at the evolutionary end points, we performed additional social DDM analyses with fixed parameter settings (i.e., no evolution of parameters). To investigate the effect of the bias z_p , boundary separation θ , and social drift s , on individuals' performance (i.e., their payoff and their hit and correct rejection rates), we varied the parameter of interest—for different group sizes and error costs—




Figure 5.2. Example trajectories of the evolutionary algorithm. The evolution of the (a) bias, (b) boundary separation, and (c) social drift rate over 1,000 generations in cooperative groups of five individuals and a miss/false alarm cost ratio of 4. The colored lines represent the average parameter values within each of the eight evolving populations and the black line indicates the average across all populations.

while fixing the other two parameters at their evolved level of cooperative groups, and measured individuals' performance over 1,000,000 repetitions (see Figs. 5.4 & 5.5a–c).

Evaluating the influence of competition. Since the end points of cooperative and competitive groups differed (see Results), we studied how competition drives populations away from the optimal end points of cooperative groups. As before, we varied the bias z_p , boundary separation θ , or social drift s while fixing the remaining two parameters at their evolved level of cooperative groups. However, we also introduced individual heterogeneity by assigning half of the individuals of each group a higher parameter value, and the other half a lower value (splitting groups of five randomly), using a difference of 0.4 for bias, 0.2 for boundary separation, and 0.1 for the social drift. This allowed us to measure the benefits of having a higher or lower parameter value than the other group members and, thereby, the effect of competition. We measured payoffs over 1,000,000 repetitions for each parameter combination (see Fig. 5.5d–f). The code for the analyses can be accessed at https://github.com/alantump/socialDDM_Evolutionary_Algorithm.

Results

In all scenarios, populations converged to a single solution, independent of the starting conditions. This indicates that the algorithm found robust solutions across group size, cost asymmetry, and competitiveness. Figure 5.2 shows examples of the trajectories of three evolving parameters for one scenario (see Supplementary Fig. D1 for other scenarios). Figure 5.3 shows the evolved parameters (i.e., the end points of the evolutionary trajectories) for different error costs and group




Figure 5.3. Outcomes of the evolutionary algorithms per group size and error cost in cooperative groups. Error cost ratio is symmetrical (1), moderately asymmetrical (2), or strongly asymmetrical (4). (a) When costs are symmetrical, no bias evolves at any group size. When costs become more asymmetrical, small groups evolve a bias; large groups do not. (b) In larger groups a higher boundary separation evolves, independent of the error costs. (c) Across all combinations of group size and error cost, a high social drift rate evolves. The dots and error bars represent the mean and standard deviation across the eight populations.

sizes in cooperative groups.

Asymmetry in costs. When both errors were equally costly (i.e., symmetrical error costs), we observed, as expected, no development of a bias in cooperative groups at any group size (Fig. 5.3a). However, as asymmetry in costs increased, individuals alone and in small cooperative groups evolved a bias towards the signal decision boundary avoiding the costly error. This was not the case in large cooperative groups. How can this effect be understood? Figure 5.4 shows the hit and correct rejection rates, as well as the payoff, for different group sizes using different biases (while fixing the boundary separation and social drift rate at the evolved level of the specific group size; see Fig. 5.3). Across all group sizes, increasing the bias increased the hit rate but decreased the correct rejection rate. Looking at the associated payoffs, we observed again that under symmetrical costs, the highest payoffs were obtained at a bias level of 0, which maximizes the combined sum of the hit and correct rejection rate. However, when misses became more costly than false alarms, single individuals and small groups maximized their payoff at a relatively high bias to avoid costly misses. In large groups, by contrast, a bias close to 0 was optimal. Large groups achieved high hit and correct rejection rates without a bias, and were very sensitive to small biases. Increasing their bias did increase their hit rate, but this did not outweigh the associated costs of the steep drop in the correct rejection rate. The steepness of this drop increased with group size. In other words, a strong bias in large groups would lead to many false alarms, which can be avoided by reducing the bias; small groups, however, cannot avoid false alarms in the presence of strong asymmetrical




Figure 5.4. The hit and correct rejection rates (black lines, left axis) and the payoff (colored lines, right axis) as a function of the bias for different group sizes and error costs for cooperative groups. Across all group sizes, increasing the bias towards the decision boundary of the signal leads to an increase in the hit rate at the expense of the correct rejection rate. Under symmetrical error costs, individuals across all group sizes maximize their payoff by maximizing the hit and correct rejection rate alike; this occurs at a bias close to 0. Under asymmetrical costs, individuals need to ensure a high hit rate in order to avoid costly misses. Small groups achieve this by developing a bias. Large groups achieve a high hit (and correct rejection) rate without a bias, and therefore maximize the payoff at a much lower bias. The boundary separation and social drift were fixed at the end points of the evolutionary algorithms for each combination of group size and error cost.

costs. Figure 5.5a further illustrates this by showing the payoffs associated with different levels of bias under highly asymmetric costs. Here, small groups can even outperform large groups under high levels of bias emphasizing the sensitivity of large groups to biases.

Boundary separation. Figure 5.3b shows that individuals in larger groups evolved a larger boundary separation and therefore waited longer for social information. This effect is independent of the asymmetry in error costs. Why should individuals wait longer in larger groups? Figure 5.5b shows the payoffs at different levels of boundary separation for different group sizes. Because the potential benefits of social information are higher in larger groups, the relative benefits of waiting for social information are also expected to be higher. Larger groups thus evolved a greater boundary separation, as shown by a shifting payoff peak to higher boundary separations for larger groups (Fig. 5.5b).

Social drift. Across all group sizes and error costs, the social drift evolved to the maximum level (Fig. 5.3c). The evolution of these extreme parameters indicates the effectiveness of a sim-

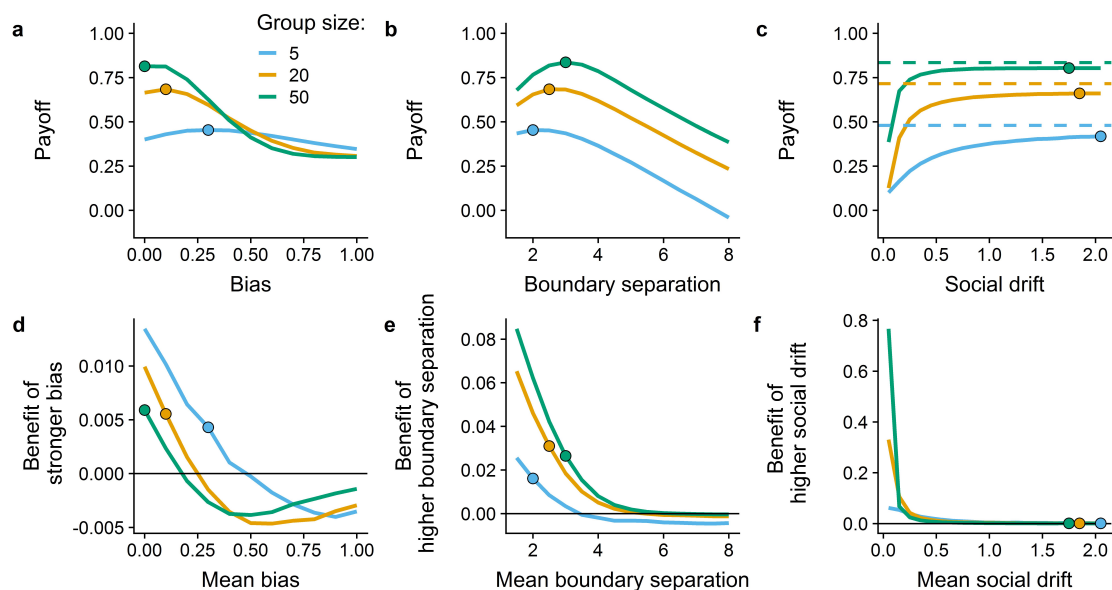


Figure 5.5. (a–c) The mean payoff of individuals in different-sized, cooperative groups across the three key parameters under high asymmetrical costs (cost asymmetry: 4). In these simulations, one of the evolved parameters was varied (x-axis), while the other two parameters were fixed at their evolved level of cooperative groups. Larger groups maximized their payoffs (indicated by circles) at (a) a lower bias and (b) higher boundary separation compared to small groups. (c) All group sizes maximized their payoff at the highest level of social drift. Dashed horizontal lines show the mean payoff of the first responder. With increasing social drift rate, the mean payoff of all group members approximates the payoff of the first responder. (d–f) The benefits of individuals in competitive groups having above-average values in the groups having above-average values in the three key parameters under high asymmetrical costs. Positive (negative) y-values indicate that individuals with above-average (below-average) values in the respective parameter achieve a higher payoff. Competitive groups are expected to evolve parameter values at which their members do not profit from having a higher—or lower—parameter value (i.e., where the colored lines meet the horizontal solid line at zero). These values partly differed from optimal outcomes in cooperative groups (circles). Individuals in competitive groups benefited from having a (d) higher bias and (e) higher boundary separation than individuals in cooperative groups. (f) Cooperative and competitive groups did not differ, however, in their optimal value of social drift.

ple “copy-the-first” heuristic, whereby individuals immediately imitate the decision of the first responder via a strong social drift. The reason this simple heuristic performs so well lies in the way personal information is gathered. Individuals with more accurate personal information start, on average, closer to a decision boundary than do individuals with less accurate information. This gives rise to a process of self-organization, with more accurate individuals making faster decisions (Tump et al., 2019). The first responder therefore generally achieves a higher payoff than do later-deciding individuals, independent of social drift rate or group size (as indicated by the higher dashed lines compared to the solid lines in Fig. 5.5c). For later-deciding individuals, the best strategy is thus to increase the social drift rate in order to imitate the first decision, thereby saving costly time.

Competitive versus cooperative groups. Next, we compared the evolutionary outcomes of cooperative and competitive groups. Across all group sizes, competitive groups developed a stronger bias towards the signal boundary than did cooperative groups (Fig. 5.6a). To investigate why, we introduced inter-individual heterogeneity in the bias level within competitive groups, and compared the payoffs of these different bias levels. We observed that, at the bias level that maximized the mean group payoff (dots in Fig. 5.5a, d), competitive individuals with a slightly higher bias gained higher individual payoffs—this advantage only disappeared when the group had a substantially higher mean bias; Fig. 5.5d). This could be due to the tension between providing good information to group members and maximizing one’s own payoff. Individuals provide better social information by reducing their bias. Increasing this bias helps an individual avoid the high costs of misses but will result in less accuracy and therefore more misleading social information.

When comparing the optimal boundary separation for cooperative and competitive groups, we found that, for all group sizes, competitive groups evolved higher boundary separations than did cooperative groups (Fig. 5.6b). Figure 5.5e shows that, at the maximum payoff level of cooperative groups (dots), competitive individuals benefited from having a slightly higher boundary separation. Over a large parameter range, individuals profited from waiting slightly longer than others for social information driving the boundary separation in competitive groups to higher values. Notably, the larger bias and boundary separation in competitive groups partly undermines the benefits of collective decision making: The mean payoff of cooperative groups increased much stronger with increasing group size compared to the payoff of competitive groups (Fig. 5.6d). Finally, both cooperative and competitive groups evolved to the maximum level of social drift (Fig. 5.6c). In line with this, individuals with a lower drift rate never outperformed individuals with a higher drift rate (indicated by the strictly positive values in Fig. 5.5f). This confirms the strong adaptive benefits of using social information, or even copying the first responder, independent of group size or cooperative setting.

Discussion

We investigated the evolution of individuals’ optimal decision rules across different group sizes, cost asymmetries, and competitiveness. When confronted with a high cost of misses (compared to the cost of false alarms), individuals in small groups evolved a bias to avoid them. This agrees with earlier findings from signal detection theory (Maddox, 2002) and has also been found as a




Figure 5.6. Comparing the evolutionary outcomes of cooperative and competitive groups at an error cost ratio of 4. The dots and error bars represent the mean and variance of the end points of the evolutionary simulations. Across all group sizes, competitive groups evolved (a) a larger bias and (b) larger boundary separation, indicating a conflict between individual- and group-level interests. (c) Both cooperative and competitive groups evolved maximum social drift rates. (d) At large, but not small, group sizes, cooperative groups outperformed competitive groups.

start point bias in the DDM framework (Mulder et al., 2012). Strong biases are adaptive in small groups, but in large groups they amplify quickly, leading groups to almost always decide for signal. This leads to a high hit rate, but also to a high false alarm rate. This implies that individuals that do not adjust their biases accordingly can be highly detrimental for large groups.

Individual response biases can have dramatic consequences for collective systems—for example, in panicking crowds. Following terrorist attacks like 9/11 or the Paris attacks in 2015, it is conceivable that some individuals adjusted their response bias towards an alarm response. This type of adjustment might be wise in a small group, but can quickly escalate in large crowds, which are more vulnerable to false alarms. In post-9/11 Chicago, for example, several club visitors mistook pepper spray for a poison gas attack; the resulting panic left 21 dead (CNN, 2003). The risk of amplification is further worsened in competitive groups, which evolve a higher bias than cooperative groups. When individuals aim to maximize their own payoff, they are willing to accept a higher level of false alarms at the expense of the collective well-being. Indeed, large competitive groups performed substantially worse than large cooperative groups, partly due to a higher evolved bias. In high-stress situations, a shift from cooperative to competitive behavior is frequently observed (Mintz, 1951; Moussaïd and Trauernicht, 2016). Taken together, our results highlight the potential danger posed to groups by individuals with a high response bias, and call for a more detailed understanding of how information spreads in such situations.

An important reason that groups are vulnerable to information cascades is their high reliance on social information. We found that social drift rate evolves to a maximum across group size, cost asymmetry, and competitiveness, indicating the superiority of a copy-the-first strategy. This finding confirms previous studies on strategic delay, which describe a Nash equilibrium in which

everyone initially delays their choice. As time passes, the individual with the best information can assume that their information is better than that of other members in the group, inferred by the absence of choice. Since the individual with the best information is expected to decide first, others then imitate this decision (Gul and Lundholm, 1995; Zhang, 1997). Adopting a copy-the-first strategy not only allows individuals to rely on the social source with the strongest evidence but also saves them time, since they do not need to wait for other individuals to decide.

Crucially, these studies—and ours—assume that individuals have the same speed–accuracy trade-off (e.g., the same boundary separation). If individuals differed in their speed–accuracy trade-off, those with the best information would be less likely to be the first responders. Indeed, in this scenario the first responder is more likely to be an individual with a high emphasis on speed—which comes at the expense of accuracy. When individuals differ in their speed–accuracy trade-off, favoring the decisions of individuals who respond quickly might undermine the benefits of social information use and of a copy-the-first strategy. Individual differences in speed–accuracy trade-offs may therefore be a critical factor in explaining why few, if any, empirical studies find that individuals simply copy the first decision maker (e.g., Kurvers et al., 2015). Future work could investigate the consequences of individual heterogeneity in speed–accuracy trade-offs on optimal decision making in collective systems.

Another interesting extension of our framework is the evolution of more complex strategies to integrate social information. For the sake of simplicity, we assumed a linear relationship of majority size and drift rate. However, individuals could also use nonlinear responses such as quorum thresholds (Kurvers et al., 2014; Marshall et al., 2019; Sumpter, 2006; Sumpter and Pratt, 2009). Nonlinear response strategies down-weight small minorities, but once the majority reaches a certain threshold they ramp up social information use.

The willingness to wait for further social information—described by the boundary separation—also influences the collective dynamics. We found that the boundary separation increased with group size, meaning that individuals in larger groups required more evidence to make a decision. This effect was particularly prominent in competitive groups where individuals profited from requiring more evidence than others. These results resemble a well-known finding in social psychology: the bystander effect. According to the bystander effect—first demonstrated by Darley and Latané (1968)—people are less likely to offer help the more other people there are nearby. We show that waiting longer to see whether others respond can be an adaptive strategy, as individuals in larger groups should only make a choice (e.g., whether to offer help) with strong

evidence. Matching this prediction, a study using CCTV footage found that increased bystander presence reduced individuals' likelihood of intervening (e.g., via increased boundary separations) while simultaneously increasing the likelihood of someone intervening (Philpot et al., 2019). The bystander effect could be explained as a rational adaptation to maximize informational gain in varying group sizes. However, future work should investigate this explanation in situations with less moral connotations.

In conclusion, we show that in the presence of asymmetrical costs, individuals should adjust their response bias to the group size to maximize their payoff. In particular, individuals in large groups should avoid strong response biases as they frequently trigger false information cascades. Further, we show that in competitive groups, individuals are—to some extent—indifferent to the potentially negative consequences of their response bias; this leads groups—especially large groups—to fail to reap the full potential of making decisions with others. As asymmetrical costs are the rule rather than the exception, our results have important implications for understanding a wide range of social dynamics, including police officers' decisions to shoot (Pleskac et al., 2018), panicking behavior in crowds (Moussaïd et al., 2016), and escape responses under predation risk (Lima, 1995).

References

- Anderson, L. R. and Holt, C. A. (1997). Information cascades in the laboratory. *The American Economic Review*, 87(5):847–862.
- Banerjee, A. V. (1992). A simple model of herd behavior. *The Quarterly Journal of Economics*, 107(3):797–817.
- Beauchamp, G. and Ruxton, G. D. (2007). False alarms and the evolution of antipredator vigilance. *Animal Behaviour*, 74(5):1199–1206.
- Ben-Yashar, R. C. and Nitzan, S. I. (1997). The optimal decision rule for fixed-size committees in dichotomous choice situations: The general result. *International Economic Review*, 38(1):175–186.
- Bikhchandani, S., Hirshleifer, D., and Welch, I. (1998). Learning from the Behavior of Others: Conformity, Fads, and Informational Cascades. *Journal of Economic Perspectives*, 12(3):151–170.
- Chittka, L., Skorupski, P., and Raine, N. E. (2009). Speed–accuracy tradeoffs in animal decision making. *Trends in Ecology & Evolution*, 24(7):400–407.
- CNN (2003). *Judge blocks charges against E2 owners*.
<https://edition.cnn.com/2003/US/Midwest/02/18/chicago.nightclub>.
- Couzin, I. D., Krause, J., Franks, N. R., and Levin, S. A. (2005). Effective leadership and decision-making in animal groups on the move. *Nature*, 433(7025):513.
- Darley, J. M. and Latané, B. (1968). Bystander intervention in emergencies: Diffusion of responsibility. *Journal of Personality and Social Psychology*, 8(4):377–383.
- Gallup, A. C., Hale, J. J., Sumpter, D. J., Garnier, S., Kacelnik, A., Krebs, J. R., and Couzin, I. D. (2012). Visual attention and the acquisition of information in human crowds. *Proceedings of the National Academy of Sciences*, 109(19):7245–7250.
- Gold, J. I. and Shadlen, M. N. (2007). The neural basis of decision making. *Annual Review of Neuroscience*, 30:535–574.

- Green, D. M. and Swets, J. A. (1966). *Signal detection theory and psychophysics*. John Wiley, New York, NY.
- Gul, F. and Lundholm, R. (1995). Endogenous timing and the clustering of agents' decisions. *Journal of Political Economy*, 103(5):1039–1066.
- Hamblin, S. (2013). On the practical usage of genetic algorithms in ecology and evolution. *Methods in Ecology and Evolution*, 4(2):184–194.
- Johnson, D. D. P., Blumstein, D. T., Fowler, J. H., and Haselton, M. G. (2013). The evolution of error: Error management, cognitive constraints, and adaptive decision-making biases. *Trends in Ecology and Evolution*, 28(8):474–481.
- Kurvers, R. H., Wolf, M., Naguib, M., and Krause, J. (2015). Self-organized flexible leadership promotes collective intelligence in human groups. *Royal Society Open Science*, 2(12):150222.
- Kurvers, R. H. J. M., Wolf, M., and Krause, J. (2014). Humans use social information to adjust their quorum thresholds adaptively in a simulated predator detection experiment. *Behavioral Ecology and Sociobiology*, 68(3):449–456.
- Lima, S. L. (1995). Collective detection of predatory attack by social foragers: fraught with ambiguity? *Animal Behaviour*, 50(4):1097–1108.
- Macmillan, N. A. and Creelman, C. D. (2005). *Detection theory: A user's guide*. Lawrence Erlbaum Associates, Mahwah, NJ.
- Maddox, W. T. (2002). Toward a unified theory of decision criterion learning in perceptual categorization. *Journal of the Experimental Analysis of Behavior*, 78(3):567–595.
- Marshall, J. A. R., Kurvers, R. H. J. M., Krause, J., and Wolf, M. (2019). Quorums enable optimal pooling of independent judgements in biological systems. *eLife*, 8:e40368.
- Mintz, A. (1951). Non-adaptive group behavior. *The Journal of Abnormal and Social Psychology*, 46(2):150–159.
- Moussaïd, M., Kapadia, M., Thrash, T., Sumner, R. W., Gross, M., Helbing, D., and Hölscher, C. (2016). Crowd behaviour during high-stress evacuations in an immersive virtual environment. *Journal of The Royal Society Interface*, 13(122):20160414.
- Moussaïd, M. and Trauernicht, M. (2016). Patterns of cooperation during collective emergencies in the help-or-escape social dilemma. *Scientific Reports*, 6:33417.
- Mulder, M. J., Wagenmakers, E.-J., Ratcliff, R., Boekel, W., and 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.
- Philpot, R., Liebst, L. S., Levine, M., Bernasco, W., and Lindegaard, M. R. (2019). Would I be helped? Cross-national CCTV footage shows that intervention is the norm in public conflicts. *American Psychologist*.
- Pleskac, T. J. and Busemeyer, J. R. (2010). Two-stage dynamic signal detection: A theory of choice, decision time, and confidence. *Psychological Review*, 117(3):864–901.
- Pleskac, T. J., Cesario, J., and Johnson, D. J. (2018). How race affects evidence accumulation during the decision to shoot. *Psychonomic Bulletin & Review*, 25(4):1301–1330.
- Ratcliff, R. and McKoon, G. (2008). The diffusion decision model: theory and data for two-choice decision tasks. *Neural Computation*, 20(4):873–922.
- Reebs, S. G. (2000). Can a minority of informed leaders determine the foraging movements of a fish shoal? *Animal Behaviour*, 59(2):403–409.
- Stroeymeyt, N., Franks, N. R., and Giurfa, M. (2011). Knowledgeable individuals lead collective decisions in ants. *Journal of Experimental Biology*, 214(18):3046–3054.
- Sumpter, D. J. T. (2006). The principles of collective animal behaviour. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 361(1465):5–22.
- Sumpter, D. J. T. and Pratt, S. C. (2009). Quorum responses and consensus decision making. *Philosophical transactions of the Royal Society of London. Series B, Biological sciences*, 364(1518):743–753.
- Swets, J. A., Dawes, R. M., and Monahan, J. (2000). Psychological science can improve diagnostic decisions. *Psychological Science in the Public Interest*, 1(1):1–26.
- Tump, A. N., Pleskac, T. J., and Kurvers, R. (2019). Wise or mad crowds? The cognitive mechanisms underlying information cascades. *PsyArXiv*.

- Watts, I., Nagy, M., Burt de Perera, T., and Biro, D. (2016). Misinformed leaders lose influence over pigeon flocks. *Biology Letters*, 12(9):20160544.
- Wolf, M., Kurvers, R. H., Ward, A. J., Krause, S., and Krause, J. (2013). Accurate decisions in an uncertain world: collective cognition increases true positives while decreasing false positives. *Proceedings of the Royal Society B: Biological Sciences*, 280(1756):20122777.
- Zhang, J. (1997). Strategic delay and the onset of investment cascades. *The RAND Journal of Economics*, 28(1):188-205.

Chapter 6

Synthesis and Future Directions

Efficiently learning from others is a complex task and requires the implementation of various social learning strategies which, in turn, crucially affect the collective dynamic. Although these dynamics often further boost the accuracy of decisions, these dynamics can also go utterly wrong. The ubiquity of social influence in our society makes it crucial to understand the emergence of beneficial or detrimental dynamics.

This thesis aimed to deepen the knowledge of social dynamics by extending several well-known models of individual decision-making into a social context. These models allowed me to understand collective dynamic using a bottom-up approach. This chapter will summarise the major findings (see also Table 6.1), put them into a broader context, and, finally, discuss future research directions.

Synthesis

Scientists from a variety of disciplines have studied what characterises good social information. One evident factor predicting information quality is the expertise of the source. However, many researchers pointed out that, in addition to the ability of these sources, their diversity is a crucial element influencing the ability to benefit from others (Sorkin et al., 2001; Luan et al., 2012; Krause et al., 2011; Hong and Page, 2004). Only a diverse group produces diverse sets of solutions which allow social learning rules operating with majorities or averaging to perform well (Herzog et al., 2019). But which features of the decision-making process contribute to the diversity of solutions in a group? Describing the cognitive process from the perspective of Brunswik's lens model (Brunswik, 1952) allows to identify features by which decision makers may come to diverse conclusions: (1) by observing different cues, (2) via diverse inference strategies, and (3) by developing different beliefs about the validity of cues. While the influence of observed cues and inference strategies on behavioural diversity has been investigated (Sorkin et al., 2001; Fujisaki et al., 2018), the potential of belief diversity is less understood. Chapter 2 addressed the role of belief diversity in a paradigm in which the meaning of cues had to be learned. Despite provided with the same information, individuals developed different cue beliefs, including many wrong ones. The resulting diversity in cue beliefs, indeed, increased social information quality. However, individuals failed to realize the full potential of diversity. Using simulations, I show that individuals could have benefited from diversity if they would had relied on simple majorities. Individuals, however, showed great reluctance to rely on social information and followed the majority only if these majorities were very large (i.e., if the majority showed strong agreement). In many situations, individuals might benefit from relying on consistent social information, as agreement can predict expertise (Kurvers

et al., 2019; Ravazzolo and Røisland, 2011). These strategies are, however, expected to strongly impair the benefit of diversity because diversity can easily be misinterpreted as disagreement and, hence, the absence of evidence.

Despite the general consent that diversity is an important ingredient to benefit from others, this chapter has shown that we need to consider the social learning strategies to infer the realized potential of diversity in social environments. Chapter 3 addressed this shortcoming and provides a better understanding of individual-level learning strategies by analysing how decision makers cope with disagreement. For this purpose, I used a cognitive toolbox approach allowing the quantification of intra- and inter-individual differences in strategies individuals use to adjust their judgment after observing others' judgements. The toolbox combines Bayesian inference with more heuristic choice rules. I find that individuals put more weight on opinions which are in close proximity to other opinions (i.e., group members agree). This strategy resembles the finding in Chapter 2 that only majorities showing high agreement are being followed. Both chapters show that individuals use the agreement among other group members as an important cue when using social information. Moreover, I found that individuals weighted others' estimates much more heavily if these estimates confirmed their own opinion. Some individuals even used a keep-heuristic (see also Moussaïd et al., 2013; Jayles et al., 2017), a strategy in which they simply ignored social information, in particular when another person agreed with the own opinion. For both strategies, social information is strongly influenced by an individual's personal opinion and the notion of being confirmed, highlighting the importance of confirmation biases in the repertoire of strategies individuals possess (Koriat et al., 1980; Nickerson, 1998; Schulz-Hardt et al., 2000).

Such confirmation biases can have a substantial effect on the collective dynamic. Using simulations, I showed that when provided with polarised information (i.e., two groups with disagreeing estimates), the individual's adjustment is predicted to be very small if someone confirms his or her own opinion. In other words, trying to convince individuals by exposing them to different views or beliefs (e.g., from outside one's filter bubble) might hardly induce any opinion shift. I further show that a confirmation bias can even cause further polarisation and promote more extreme views. These results highlight the importance of understanding the social learning strategies individuals use to predict opinion dynamics (Lorenz, 2007) and to counteract the spread of misinformation—a matter of great public concern (Lewandowsky et al., 2017; Vosoughi et al., 2018).

Using a cognitive toolbox approach thereby allows testing various plausible strategies by accounting for personal information and social information but can also consider the consistency of

the individuals' behaviour over trials by having a theoretical basis for the sources of stochasticity. However, the notion of such strategy toolboxes has also received criticism. One argument is the difficulty of falsifying such models and testing them against alternative models of cognition (Todd and Gigerenzer, 2001). The main reason is its flexibility, as strategy sets in toolbox models can, in principle, contain an infinite number of strategies. Therefore, I restricted the number of investigated strategies to a relatively small set. Yet, future work should investigate further plausible strategies. This model represents a theoretical and methodological advancement and lays an important foundation for future researchers to study how individuals adjust estimates or opinions when confronted with the beliefs of others. Only a deep understanding of how such local strategies eventually generate collective patterns of opinion change allows us to meet the challenges of our globalized and interconnected world.

Using the toolbox, I investigated learning strategies in a controlled and static social environment, allowing measurement of information transmission and enabling conjectures about which collective dynamics arise. These resulting dynamics are, however, likely to result in a feedback loop that can alter the use of learning strategies. Hence, when analysing how individuals navigate in social environments, I might miss a crucial aspect when only focusing on learning strategies under static conditions. In Chapter 4, we embed individuals in a temporal dynamic system allowing for the comprehensive investigation of social learning strategies in interaction with the unfolding dynamic.

Across a wide range of situations such as consumer choice (Chen, 2008) or decisions about crossing the street (Faria et al., 2010; Pfeffer and Hunter, 2013), individuals do not decide simultaneous but sequentially. They can strategically time their decision, allowing them to decide early or to wait for (further) social information. The social DDM introduced in Chapter 4 is well-suited to investigate such a decision process and additionally to capture the resulting dynamic by decomposing choices and response times into psychologically meaningful parameters. Testing the model in an empirical study, I showed that enabling individuals to time their decision gave rise to a collective dynamic where accurate individuals tended to decide early and thereby provided correct information to others. Later-deciding individuals, in turn, were likely to follow those who already decided, which further amplified the high-quality signal. This self-organization according to information quality relies on two driving factors. First, individuals need to have a valid sense of their information quality (i.e., confident individuals are more accurate). Second, more confident individuals need to decide earlier. Importantly, the social DDM was able to recover these dynamics

and revealed the cognitive underpinnings.

The quality of information cascades is crucially influenced by the accuracy of early-deciding individuals. As only an understanding of cognitive process allows one to account for decision timing (i.e., fast vs. slow errors), the social DDM is a helpful tool to better predict the spread of accurate or inaccurate information. Hence, the social DDM makes it possible to analyse how individuals navigate in social environments and thereby predict the resulting dynamic whenever publicly observable choices are made sequentially. This framework allows future research to extend the social DDM to other areas such as crowd dynamics (Moussaïd et al., 2016), the spread of information in social media (Vosoughi et al., 2018) or collective behaviour in animals (Danchin et al., 2004).

The beneficial collective dynamic identified here is an illustrative example of how individuals coordinate without the intention to benefit the “collective well-being.” Importantly, this dynamic predominately relies on the individual’s adjustment of the decision time according to the quality of their personal information. If individuals coordinate in such a way that more accurate individuals decide early and less accurate ones decide late, this will result in a win-win situations. Hence, individuals have an incentive to accurately evaluate information quality and time their decision accordingly. This makes the information exchange robust to dishonest behaviour because the sender and receiver of social information each profit from sensible use of their confidence. Other studies have shown that individuals can also directly use shared confidences to weight others’ judgments (Moussaïd et al., 2013; Koriat, 2012; Zarnoth and Sniezek, 1997). Here, providing an honest confidence judgment might not be in one’s own interest. They, for example, might be neutral regarding whether others perform well or try to gain influence by providing dishonest high confidence ratings (Hertz et al., 2017). Future work should investigate under which conditions decision time might be a more accurate predictor of information quality than confidence ratings.

After having identified the general principles of how individuals incorporate personal and social information over time, I used these insights to investigate how individuals should optimally integrate information and time their decision. In sequential decision-making contexts, social interactions have a strategic component because the optimal behaviour of individuals strongly depends on the behaviour of others—i.e., this social situation represents game-theoretical problems. In Chapter 5, I applied evolutionary algorithms, to identify advantageous strategies individuals are expected to adopt when facing a game-theoretical problem. I found that the time individuals take to make a decision is expected to depend on the quality available to them. Interestingly, individ-

uals should develop a strategy where all imitate the first responder. The intuition behind such a “copy-the-first” heuristic is that the individual with the most reliable evidence is likely to decide first. The later-deciding individuals will, on average, have less evidence and mimic the first one (Gul and Lundholm, 1995; Zhang, 1997; Bikhchandani et al., 1998). Although the whole group outcome is based on a single decision, this strategy performs substantially better than a single individual would do, simply because of the strategic timing of decisions.

Yet, such strong social information use comes with a cost. When confronted with asymmetric error costs, single individuals are expected to develop response biases to avoid the more costly outcome. Though evolutionarily advantageous in small groups, response biases in large groups are likely to get amplified and trigger false information cascades. As a consequence, individuals in larger groups should develop only very weak response biases, allowing them to benefit from social information. These results highlight the danger of persistent response biases coupled with strong information use in collective contexts like panicking crowds (Moussaïd et al., 2016) or escape responses (Lima, 1995), as biased individuals are likely to trigger the spread of false alarms. In conclusion, individual-level strategies crucially influence the spread of information. Individuals are expected to self-organize by timing their decision according to their information quality and thereby boost the spread of correct information. Yet, I also show that these dynamics are sensitive to how people process the information they receive; i.e., how they derive levels of confidence and develop response biases. The social DDM allows a new perspective on such dynamics and provides a novel approach to designing and analysing experiments. In the next section, I will present future research directions in collective behaviour.

Future Research Directions

From whom to learn

Across all chapters, participants (or simulated agents) were fully interconnected and received anonymous information, limiting the possibility of strategically learning from specific others. However, previous work has shown that individuals’ reliance on others varies with specific characteristics of these group members such as hierarchical status, sex, age or reputation (Reebs, 2001; Deaner et al., 2005; Yaniv and Kleinberger, 2000). Such strategies could amplify the spread of accurate information if they allow the filtering out of less reliable information. On the other hand, if they lead to biased information aggregation—e.g., by relying on a less diverse subgroup— they

are expected to undermine the potential for collective intelligence (Krause et al., 2010; Page, 2008). Future work should, therefore, investigate how strategically learning from specific others influences collective dynamics (but see, Watts et al., 2016).

Another important issue for future research to examine is the influence of the underlying structure of the group. Information often flows through highly structured networks. The impact of the network's structure is determined by the interaction process and the social learning strategies individuals use. Highly connected networks, for example, promote the spread of accurate information if individuals tend to explore independently. On the contrary, if individuals rely little on social information, connected networks are expected to lead to higher collective gains (Mason and Watts, 2012; Lazer and Friedman, 2007). The social DDM would allow for gaining a deeper understanding of how decision timing and social information use mediate information flow across different types of social networks.

Influential minorities

Conflicting interests among group members are common when trying to find (democratic) consensus (Conradt and Roper, 2005). Previous work has suggested that uninformed or naive individuals can strongly influence the collective outcome. While some argue that such uninformed individuals are vulnerable to opinionated minorities (Olson, 2009), others have pointed out that they could actually return control to the majority (Couzin et al., 2011). Minority opinions such as corporate interests or political movements have a long history of seeking to influence public debate, often against the interest of the general public (Lewandowsky et al., 2017). We need to understand under what conditions opinionated minorities can sway majorities and which role uninformed individuals play in such collective dynamics. Can, for example, early-deciding “loud” minorities persuade uninformed individuals and, in turn, sway the later-deciding “silent” majority? The social DDM can provide insights by generating hypotheses and allowing to test these empirically.

Humans and Machines

Our society faces the constantly growing prevalence and importance of algorithms in our daily lives. Ranking algorithms and social media bots influence who observes which information. Autonomous cars drive on the streets and trading software autonomously purchases and sells financial products. This ubiquity of algorithms confronts researchers with new research questions.

First, as humans increasingly interact with machines, we need a better understanding of how these algorithms will influence our social interactions, social learning strategies and the resulting dynamics. Autonomous cars, for example, will inadvertently produce “social” information via their actions. This information will have different characteristics than information conveyed by their living counterparts. Autonomous cars will adapt their driving speed well before a traffic jam and anticipate crossing passengers earlier and in other situations than humans. It is, however, known that driving behaviour is strongly influenced by the behaviour of other drivers (Åberg et al., 1997). To predict how social systems will be altered by the introduction of intelligent algorithms into our lives, we need further studies of how people will adjust their learning strategies to this new environment. Second, we need to understand the behaviour of the algorithms themselves. Already, very simple algorithms can produce complex collective behaviours which sometimes even their engineers cannot explain (Bak et al., 1989; Tsvetkova et al., 2017), effectively turning them into “black boxes” (Voosen, 2017). To understand the behaviour of such algorithms, we need to analyse them using techniques from cognitive science (Rahwan et al., 2019). Third, to improve the performance of algorithms and machines, we need to understand how to best make use of others’ information— from humans or other machines. Should a trading algorithm, for example, buy products bought by other trading algorithms, and are “social” information using algorithms in danger of triggering information cascades?

For all of these questions, we need comprehensive cognitive models which allow us to understand all levels of social interactions: from accumulating personal information to the implementation of learning strategies to the unfolding dynamic. I believe investigating these questions in a theory-driven manner with the computational models presented in this thesis allows new insights in this direction.

Table 6.1. Summary of the findings.

Chapter	Interaction type	Cognitive model	Social learning strategies	Unfolding dynamic
2	Choice adjustment after receiving global social information	Lens Model	Individuals (1) were reluctant to incorporate social information but, (2) over time, got more sensitive to social information provided by others and (3) adjusted their sensitivity depending on their own performance.	(1) Individuals developed diverse beliefs in informative and uninformative cues. (2) Groups with diverse cue beliefs provided better information. Individuals (3) benefited from social learning but (4) failed to exploit the potential of diversity because of over-reliance on their initial opinion.
3	Estimate adjustment after receiving global social information	Cognitive Toolbox	(1) Individuals' adjustment strategies included a "keep" heuristic influenced by to the distance to the nearest other estimate. They weighted social information more: (2) the stronger the agreement with the own estimate (i.e., confirmation bias) and (3) the stronger the agreement among social information.	In scenarios with polarised opinions (i.e., two disagreeing groups) with a focal individual slightly leaning towards one view I find that: (1) individuals with a strong confirmation bias are expected to shift their opinion towards the closer agreeing groups (i.e., further polarisation) while (2) individuals with a weak confirmation bias move towards the global mean.
4	Sequential decisions	Social DDM	(1) The Individuals' confidence and personal belief determined their start point of the evidence accumulation process and (2) individuals reinforced their initial belief over time. (3) Individuals incorporated social information via an drift towards the choice preferred by the majority.	More confident individuals (1) were, on average, more accurate and (2) decided faster. As a result, unconfident and less accurate individuals, on average, decided later and received high quality information. Groups show, thus, a (3) beneficial self-organisation according to information quality.
5	Sequential decisions (simulations)	Social DDM	(1) Under asymmetrical error cost individuals in small, but not in large, groups developed response biases. (2) Individuals waited for more evidence, the larger the group. (3) Individuals imitate the first responder and, hence, develop a "copy-the-first" heuristic.	(1) The reliance on the "copy-the-first" heuristic triggers information cascades. (2) Groups are sensitive to response biases because even small biases are amplified. (3) Individuals in small, but not in large, groups are willing to accept the detrimental consequences of response biases to avoid costly errors.

References

- Åberg, L., Larsen, L., Glad, A., and Beilinson, L. (1997). Observed vehicle speed and drivers' perceived speed of others. *Applied Psychology*, 46(3):287–302.
- Bak, P., Chen, K., and Creutz, M. (1989). Self-organized criticality in the game of life. *Nature*, 342(6251):780 – 782.
- Bikhchandani, S., Hirshleifer, D., and Welch, I. (1998). Learning from the Behavior of Others: Conformity, Fads, and Informational Cascades. *Journal of Economic Perspectives*, 12(3):151–170.
- Brunswik, E. (1952). The conceptual framework of psychology. *Psychological Bulletin*, 49(6):654–656.
- Chen, Y.-F. (2008). Herd behavior in purchasing books online. *Computers in Human Behavior*, 24(5):1977–1992.
- Conradt, L. and Roper, T. J. (2005). Consensus decision making in animals. *Trends in Ecology & Evolution*, 20(8):449–456.
- Couzin, I. D., Ioannou, C. C., Demirel, G., Gross, T., Torney, C. J., Hartnett, A., Conradt, L., Levin, S. A., and Leonard, N. E. (2011). Uninformed individuals promote democratic consensus in animal groups. *Science*, 334(6062):1578–1580.
- Danchin, E., Giraldeau, L.-A., Valone, T. J., and Wagner, R. H. (2004). Public information: from nosy neighbors to cultural evolution. *Science*, 305(5683):487–491.
- Deaner, R. O., Khera, A. V., and Platt, M. L. (2005). Monkeys pay per view: adaptive valuation of social images by rhesus macaques. *Current Biology*, 15(6):543–548.
- Faria, J. J., Krause, S., and Krause, J. (2010). Collective behavior in road crossing pedestrians: The role of social information. *Behavioral Ecology*, 21(6):1236–1242.
- Fujisaki, I., Honda, H., and Ueda, K. (2018). Diversity of inference strategies can enhance the ‘wisdom-of-crowds’ effect. *Palgrave Communications*, 4(1):107.
- Gul, F. and Lundholm, R. (1995). Endogenous timing and the clustering of agents' decisions. *Journal of Political Economy*, 103(5):1039–1066.
- Hertz, U., Palminteri, S., Brunetti, S., Olesen, C., Frith, C. D., and Bahrami, B. (2017). Neural computations underpinning the strategic management of influence in advice giving. *Nature Communications*, 8(2191).
- Herzog, S. M., Litvinova, A., Yahosseini, K. S., Tump, A. N., and Kurvers, R. H. (2019). The ecological rationality of the wisdom of crowds. In *Taming Uncertainty*. MIT Press.
- Hong, L. and Page, S. E. (2004). Groups of diverse problem solvers can outperform groups of high-ability problem solvers. *Proceedings of the National Academy of Sciences*, 101(46):16385–16389.
- Jayles, B., Kim, H.-r., Escobedo, R., Cezera, S., Blanchet, A., Kameda, T., Sire, C., and Theraulaz, G. (2017). How social information can improve estimation accuracy in human groups. *Proceedings of the National Academy of Sciences*, 114(47):12620–12625.
- Koriat, A. (2012). When are two heads better than one and why? *Science*, 336(6079):360–362.
- Koriat, A., Lichtenstein, S., and Fischhoff, B. (1980). Reasons for confidence. *Journal of Experimental Psychology: Human Learning and Memory*, 6(2):107–118.
- Krause, J., Ruxton, G. D., and Krause, S. (2010). Swarm intelligence in animals and humans. *Trends in Ecology & Evolution*, 25(1):28–34.
- Krause, S., James, R., Faria, J. J., Ruxton, G. D., and Krause, J. (2011). Swarm intelligence in humans: diversity can trump ability. *Animal Behaviour*, 81(5):941–948.
- Kurvers, R. H. J. M., Herzog, S. M., Hertwig, R., Krause, J., Moussaid, M., Argenziano, G., Zalaudek, I., Carney, P. A., and Wolf, M. (2019). How to detect high-performing individuals and groups: Decision similarity predicts accuracy. *Science Advances*, 5(11).
- Lazer, D. and Friedman, A. (2007). The network structure of exploration and exploitation. *Administrative Science Quarterly*, 52(4):667–694.
- Lewandowsky, S., Ecker, U. K., and Cook, J. (2017). Beyond misinformation: Understanding and coping with the “post-truth” era. *Journal of Applied Research in Memory and Cognition*, 6(4):353–369.
- Lima, S. L. (1995). Collective detection of predatory attack by social foragers: fraught with ambiguity? *Animal Behaviour*, 50(4):1097–1108.
- Lorenz, J. (2007). Continuous opinion dynamics under bounded confidence: A survey. *International*

- Journal of Modern Physics C*, 18(12):1819–1838.
- Luan, S., Katsikopoulos, K. V., and Reimer, T. (2012). When does diversity trump ability (and vice versa) in group decision making? A simulation study. *PLoS ONE*, 7(2):e31043.
- Mason, W. and Watts, D. J. (2012). Collaborative learning in networks. *Proceedings of the National Academy of Sciences*, 109(3):764–769.
- Moussaïd, M., Kämmer, J. E., Analytis, P. P., and Neth, H. (2013). Social Influence and the Collective Dynamics of Opinion Formation. *PLoS ONE*, 8(11):e78433.
- Moussaïd, M., Kapadia, M., Thrash, T., Sumner, R. W., Gross, M., Helbing, D., and Hölscher, C. (2016). Crowd behaviour during high-stress evacuations in an immersive virtual environment. *Journal of The Royal Society Interface*, 13(122):20160414.
- Nickerson, R. S. (1998). Confirmation bias: A ubiquitous phenomenon in many guises. *Review of General Psychology*, 2(2):175–220.
- Olson, M. (2009). *The logic of collective action*, volume 124. Harvard University Press.
- Page, S. E. (2008). *The Difference: How the Power of Diversity Creates Better Groups, Firms, Schools, and Societies-New Edition*. Princeton University Press.
- Pfeffer, K. and Hunter, E. (2013). The effects of peer influence on adolescent pedestrian road-crossing decisions. *Traffic Injury Prevention*, 14(4):434–440.
- Rahwan, I., Cebrian, M., Obradovich, N., Bongard, J., Bonnefon, J.-F., Breazeal, C., Crandall, J. W., Christakis, N. A., Couzin, I. D., Jackson, M. O., et al. (2019). Machine behaviour. *Nature*, 568.
- Ravazzolo, F. and Røisland, Ø. (2011). Why do people place lower weight on advice far from their own initial opinion? *Economics Letters*, 112(1):63–66.
- Reebs, S. G. (2001). Influence of body size on leadership in shoals of golden shiners, *notemigonus crysoleucas*. *Behaviour*, 138(7):797–809.
- Schulz-Hardt, S., Frey, D., Lüthgens, C., and Moscovici, S. (2000). Biased information search in group decision making. *Journal of Personality and Social Psychology*, 78(4):655–669.
- Sorkin, R., Hays, C., and West, R. (2001). Signal-detection analysis of group decision making. *Psychological Review*, 108(1):183–203.
- Todd, P. M. and Gigerenzer, G. (2001). Putting naturalistic decision making into the adaptive toolbox. *Journal of Behavioral Decision Making*, 14(5):381–383.
- Tsvetkova, M., García-Gavilanes, R., Floridi, L., and Yasseri, T. (2017). Even good bots fight: The case of wikipedia. *PLoS ONE*, 12(2):1–13.
- Voosen, P. (2017). The AI detectives. *Science*, 357(6346):22–27.
- Vosoughi, S., Roy, D., and Aral, S. (2018). The spread of true and false news online. *Science*, 359(6380):1146–1151.
- Watts, I., Nagy, M., Burt de Perera, T., and Biro, D. (2016). Misinformed leaders lose influence over pigeon flocks. *Biology Letters*, 12(9):20160544.
- Yaniv, I. and Kleinberger, E. (2000). Advice taking in decision making: Egocentric discounting and reputation formation. *Organizational Behavior and Human Decision Processes*, 83(2):260 – 281.
- Zarnoth, P. and Snizek, J. A. (1997). The social influence of confidence in group decision making. *Journal of Experimental Social Psychology*, 33(4):345–366.
- Zhang, J. (1997). Strategic delay and the onset of investment cascades. *The RAND Journal of Economics*, 28(1):188–205.

Appendices

B | Supplementary Material to Chapter 3

Supplementary Figures

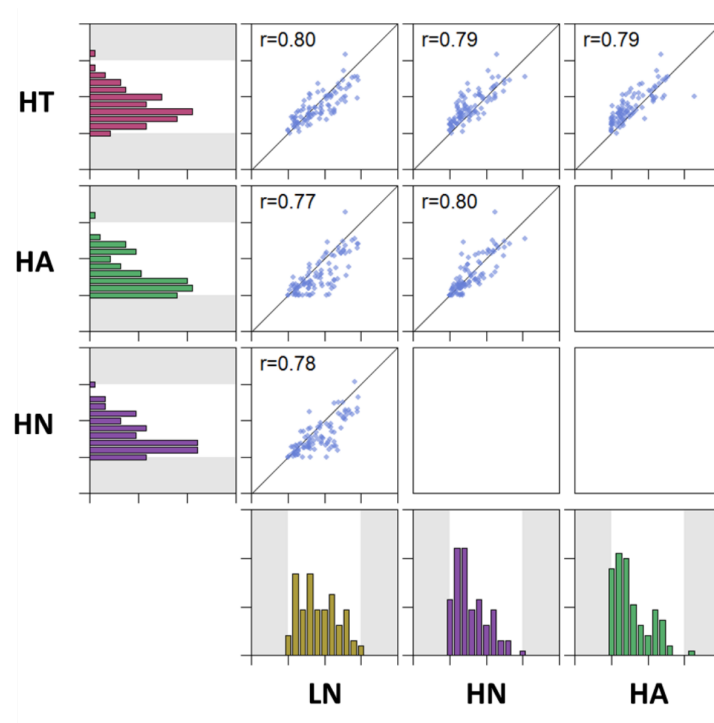


Figure B1. Distributions of individuals’ mean adjustment and correlations across the four experimental conditions. Outer histogram panels show frequency distributions of individuals’ mean adjustments (\bar{s}) towards the mean social information for each condition (LN: low variance, no skewness; HN: high variance, no skewness; HA: high variance, skewness leaning away from the participants first estimate; HT: high variance, skewness leaning towards the participant’s first estimate). Left (right) grey areas in each panel represent values of (\bar{s}) below 0 (above 1). Across conditions, almost all individuals had, on average, an (\bar{s}) value between 0 and 1, implying they did not adjusted away from the social information (i.e., $\bar{s} < 0$), nor adjusted beyond the mean social information (i.e., $\bar{s} > 1$). Scatter plots show the correlations between individuals’ mean (\bar{s}) across the respective conditions. Dots represent individuals, and shown is the Pearson’s r . Overall, we find strong correlations between participants’ mean adjustments across conditions (all Pearson correlations $P < 0.001$), indicating strong inter-individual differences in social information use. Furthermore, individuals’ mean adjustments averaged across all conditions correlated in the expected directions with self-reported questionnaire scales measuring conformity (Pearson’s $r = 0.258, CI = [0.059, 0.436]$), individualism ($r = -0.204, CI = [-0.389, -0.002]$), and resistance to peer influence ($r = -0.233, CI = [-0.415, -0.033]$), confirming previous findings within this paradigm (Molleman et al., 2019).

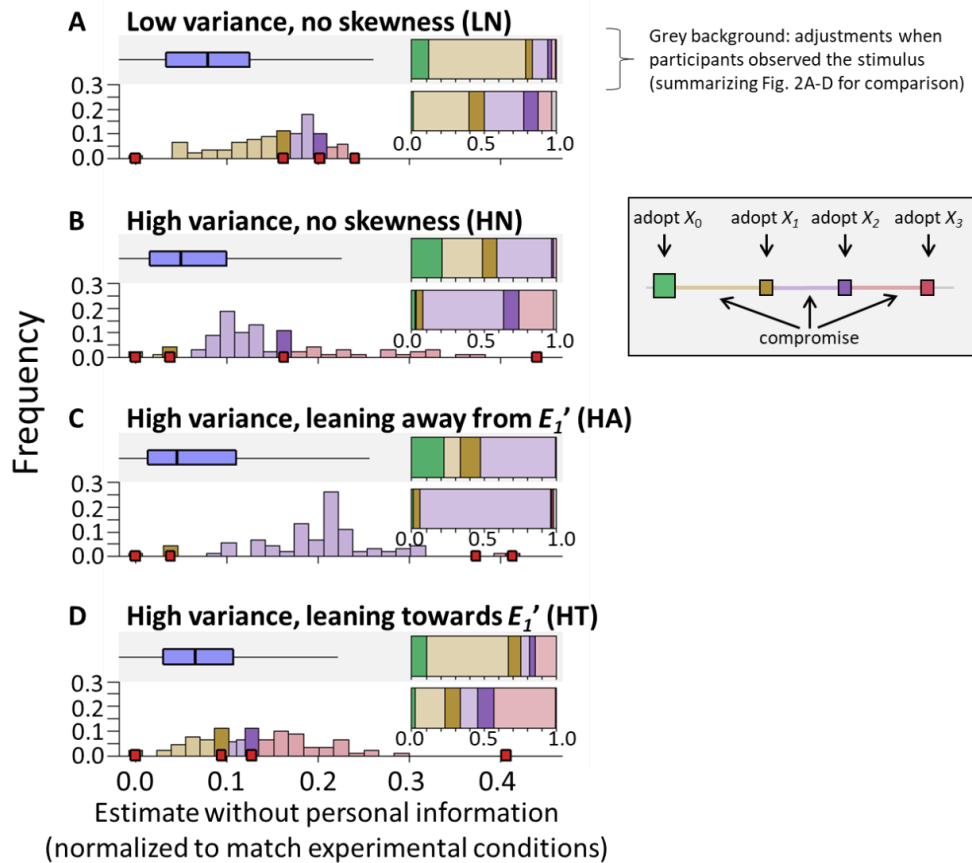


Figure B2. Adjustments in control conditions in which participants did not observe the stimulus, but only observed four peer estimates. Trials were created by first drawing one random estimate (E_1') from a pool of prerecorded estimates. Then, three additional pre-recorded estimates were drawn analogous to the four experimental conditions (see Methods). Data has been normalized in the same way as Fig. 3.2A-D, taking (E_1') as the reference point. Histograms show the distribution of estimates in the same colour coding as Fig. 3.2A-D. For comparison, in the grey background on top of each panel we summarize behavior in conditions where participants did observe the stimulus (Fig. 3.2A-D; blue boxplot: distributions of relative estimates; insets: relative frequencies of each of the qualitative cases).




Figure B3. Both in experimental (i.e., with observing the stimulus; colored dots and boxplots) and control (i.e., without stimulus; grey dots and boxplots) conditions, participants systematically deviated from an ‘ideal Bayesian observer’, who would weigh all estimates equally. Boxplots and dots show the absolute difference between participants’ second estimates and the arithmetic mean of all four estimates (see Methods for more details). Participants deviated more from the arithmetic mean in the experimental conditions compared to the control conditions, most likely because of egocentric discounting.

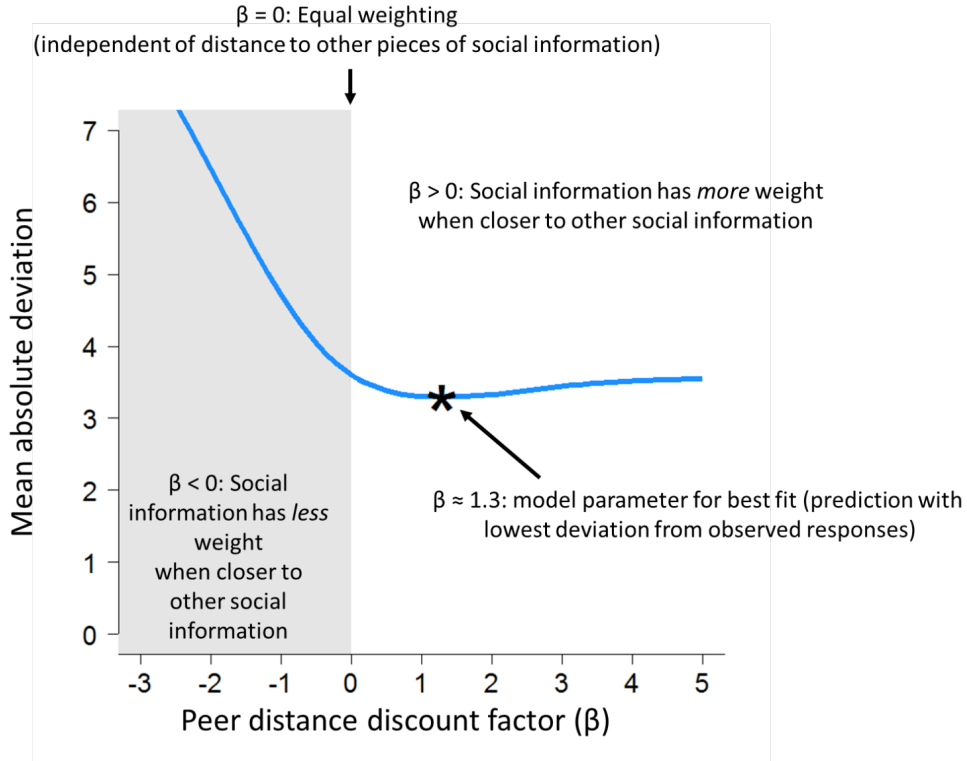


Figure B4. Descriptive model fitted to responses The discrepancy of actual and expected (E'_2) with varying proximity discounting in the control rounds in which participants did not observe the stimulus. The blue line shows mean absolute difference between per-round predictions and observed responses as a function of parameter β (capturing ‘proximity weighting’ as the extent of discounting of social information that is inconsistent with other social information). This model assumes that each response is an average of the pieces of social information (X_i), weighted according to their summed distance (d) to other pieces of social information. Formally, predicted responses are calculated as: $\widehat{E}'_2 = \sum_{j=1}^4 w_j \times X_j$. In this formula, the weighting of each piece of social information (w_i) is determined by discount factor β : $w_i = \frac{d_i^{-\beta}}{\sum_{j=1}^4 d_j^{-\beta}}$, where $d_i = \sum_{j=1}^4 |X_i - X_j|$. When $\beta > 0$, more weight is assigned to social information when it is closer to other social information. We observe that the model predictions match the observed responses best when $\beta \approx 1.3$. We interpret this as evidence of a ‘peer proximity effect’ (or a ‘consensus effect’): on average, individuals tend to assign more weight to social information when it is closer to other social information.

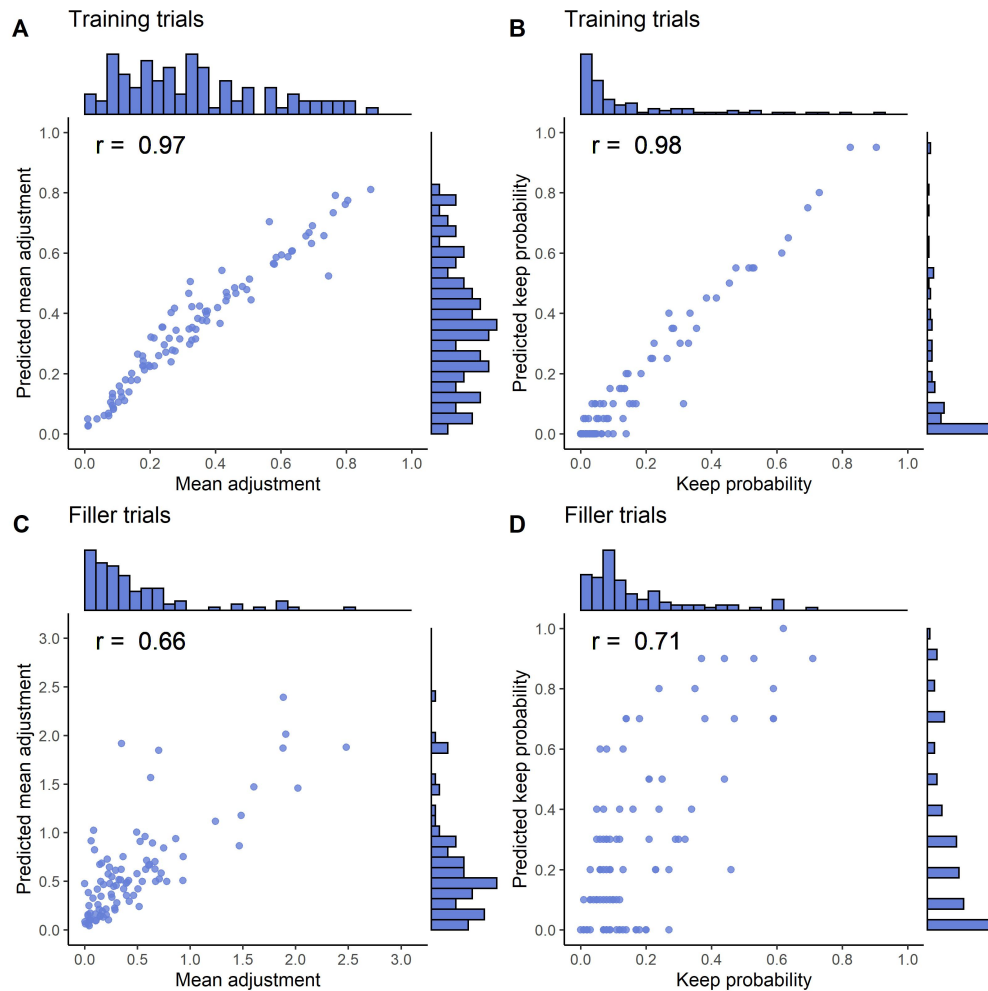


Figure B5. Correlation of predictions of the best-fitting cognitive model (y-axis) with the empirical data (x-axis). (A,B) Predictions for the 20 rounds of the 4 experimental conditions, to which the model was fitted (i.e. the ‘training set’), with (A) showing mean adjustment for each individual and (B) showing the fraction of rounds in which they chose to choose to ‘keep’ their initial estimate. We observe that the predictions of the cognitive model closely match data in the experimental conditions. (C,D) Predictions for the ‘filler’ trials with social information randomly drawn from previous participants. Again, the behaviour is fairly well predicted by the model. As for previous predictions we sampled 10 estimates for each individual and trial. In each panel, numbers in the top left corner indicate Pearson’s correlations.

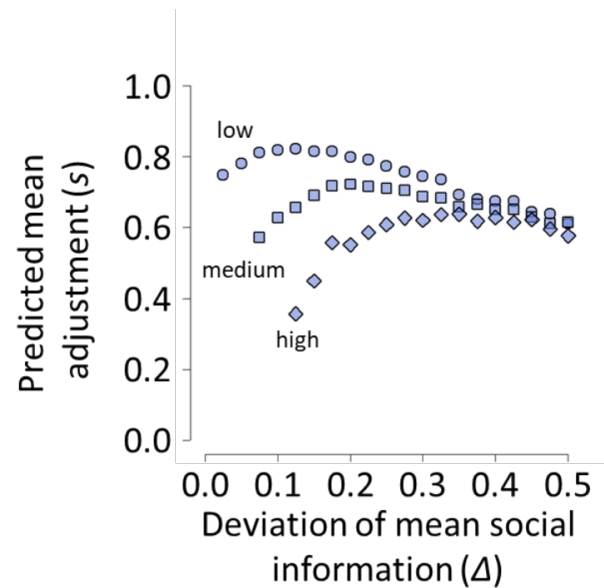


Figure B6. Social information tends to have the strongest impact when at intermediate distance. We simulated trials that systematically increasing the distance between individuals own initial estimates and the mean social information ($\Delta = (\bar{X} - E_1)/E_1$). Dots show predicted mean adjustments as a function of the mean distance to three pieces of social information, for three levels of ‘proximity’, with peer estimates at 0 (low), 3 (medium) or 6 (high) numbers away from each other. We observe that mean shifts for each level are highest when the distance of social information is intermediate, confirming previous observations (Moussaïd et al., 2013; Jayles et al., 2017). When the distance to mean social information is very low, individuals are very likely to keep their initial estimates. When the distance to mean social information is very high, individuals are likely to compromise, but assign little weight to social information, again leading to reduced average adjustments. At intermediate distance, social information has the strongest impact: in those cases, the likelihood that individuals keep their initial estimates is very close to zero, but when compromising, they still assign weight to social information. These effects hold across various levels of peer proximity. Simulations with peers farther away from each other start at higher values of Δ to ensure that all pieces of social information were in the same direction from the first estimate.

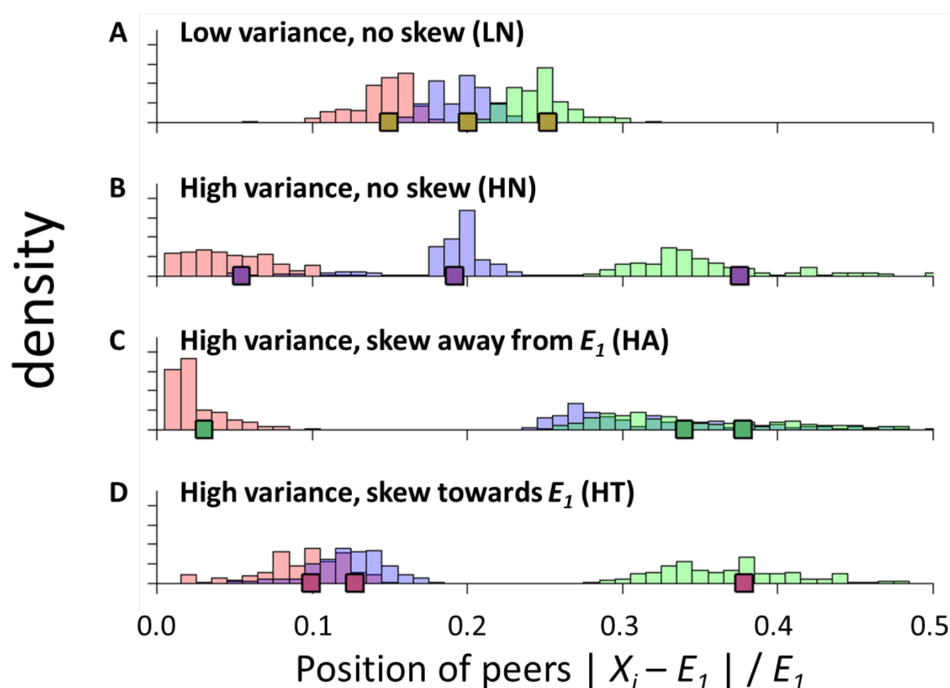


Figure B7. Relative position of peers as shown in each of the experimental conditions. Social information was based on data from a prerecorded pool of 100 MTurkers who completed the task without social information. As described in the Methods, we assigned an experimental condition to each of the 30 rounds and calculated for each value of E_1 the ‘triple’ of previous estimates that most closely matched the experimental condition (using the lowest value resulting in the ‘cost function’; see Supplementary Methods). The graph shows for each condition where the peers were located in the actual experiment, as a fraction of E_1 . Color coding: red: nearest peer, blue: middle peer, green: farthest peer. Dots show for each of these three peers their average position (Fig. 3.1E). The graph omits data points that exceeded 0.5, the vast majority of which are in the Random condition, which was used as ‘filler’ rounds. The Supplementary Methods below provide a formal description of how we defined these experimental conditions.

Supplementary Tables

Table B1. Determinants of individuals’ social information use. Numbers on the left hand side show results from a Bayesian linear mixed model fitted to individuals’ average adjustments towards the mean social information across the rounds of each of the treatments, with ‘individual’ as random effect. On the right hand side, we show pairwise comparisons between experimental conditions. Average adjustments were calculated as the relative distance adjusted towards the mean social information (Fig. 3.1D), across the five rounds of each of the four experimental conditions (LN, HN, HA, HT; Fig. 3.1E), yielding four data points for each participant. We omitted those rounds in which a participant adjusted away from the mean social information. The LN condition was used as the baseline. Relative to that baseline, the HN and HA conditions had a highly significant negative effect. The HT condition had a smaller negative effect. Pairwise comparisons indicate that, except for the HN and HA conditions, all pairs of experimental conditions differed significantly from each other. Age and gender did not significantly affect social information use. To quantify individuals’ constancy of displayed behaviour across conditions we also derived the commonly used index of ‘repeatability’ for this model ($R=0.790$ [0.697, 0.839]; Nakagawa & Schielzeth, 2010; Stoffel et al., 2017).

	Estimate [95% CI]	Pairwise comparison	Estimate [95% CI]
Low variance, No skew (LN); baseline	0.39 [0.20, 0.57]	HN - LN	-0.13 [-0.15, -0.10]
High variance, No skew (HN)	-0.13 [-0.16, -0.09]	HA - LN	-0.14 [-0.16, -0.11]
High variance, skew Away from E1 (HA)	-0.14 [-0.17, -0.10]	HT - LN	-0.05 [-0.08, -0.02]
High variance, skew Towards E1 (HT)	-0.05 [-0.08, -0.02]	HA - HN	-0.01 [-0.03, 0.02]
Age	0.00 [0.00, 0.00]	HT - HN	0.08 [0.05, 0.10]
Gender (0=male, 1=female)	0.02 [-0.07, 0.12]	HT - HA	0.09 [0.06, 0.11]
N	95		
n	380		

Table B2. Effects of experimental conditions on individuals’ use of adjustment strategies. Values indicate for each treatment the predicted frequencies of adjustment strategies - keeping, compromising, adopt the nearest peer (X_1), and all ‘other’ cases pooled - from a mixed multinomial regression with ‘participant’ as random effect (insets Fig. 3.2A-D of the main text). Values in brackets indicate 95% confidence intervals. Apart from the experimental conditions, the full model also included age and gender, neither of which had credible non-zero effects on any of the relative frequencies.

condition	keep	compromise	copy X1	other
LN	0.07 [0.04, 0.10]	0.79 [0.70, 0.86]	0.02 [0.01, 0.05]	0.12 [0.07, 0.19]
HN	0.16 [0.11, 0.22]	0.26 [0.18, 0.37]	0.07 [0.04, 0.11]	0.51 [0.37, 0.63]
HA	0.16 [0.11, 0.23]	0.08 [0.05, 0.13]	0.09 [0.05, 0.15]	0.67 [0.54, 0.77]
HT	0.06 [0.04, 0.09]	0.67 [0.56, 0.76]	0.06 [0.03, 0.10]	0.21 [0.13, 0.32]

Table B3. The comparison of all versions of the cognitive models with all possible combinations of considered features. We compared model with (1) or without (0) keeping, adopting, proximity and confirmation, and calculated the looic. The looic_diff values indicate the goodness-of-fit of each model compared to the model with the lowest looic (i.e., rank 1). The best-fitting model includes the keep heuristic, proximity and confirmation, but not the ‘adopt’ heuristic.

Keep	Adopt	Proximity	Confirmation	Looic	Looic_diff	Rank
1	0	1	1	8205	0	1
1	0	0	1	8222	-16.95	2
1	1	1	1	8228	-22.54	3
1	1	0	1	8241	-35.48	4
1	1	1	0	8424	-219.32	5
1	0	1	0	8446	-240.41	6
1	1	0	0	8463	-257.83	7
1	0	0	0	8492	-287.15	8
0	1	1	1	8687	-482.24	9
0	0	1	1	8702	-497.17	10
0	1	0	1	8716	-510.48	11
0	0	0	1	8718	-513.05	12
0	1	1	0	8964	-758.91	13
0	0	1	0	8972	-766.67	14
0	1	0	0	8999	-794.1	15
0	0	0	0	9033	-828.32	16

Table B4. The parameters shaping the group-level parent distributions, their priors, their upper (/lower) bounds and the model estimates of the best fitting model (i.e, without adopt). The parameters describe the mean and standard deviation of parent normal distributions and the two parameters shaping the of parent beta distributions. Standard deviations and parameters describing the beta distributions were truncated at 0.01 to avoid zero or negative values. Subscripts p and s respectively indicate personal and social information; the other subscripts refer to their respective model feature (distance or proximity weighting; heuristics of keeping or adopting). The right hand side column shows the parameter estimates of the best-fitting model. This model did not include the ‘adopt’ heuristic (see Table B3) so no values are shown for the parameters associated with that heuristic.

Group-level parameters	Priors	Bounds (max/min)	Parameter estimates
$\mu^{\sigma_p^2}$	N(10, 5)	1 / Inf	3.88 [3.5, 4.42]
$\tau^{\sigma_p^2}$	N(0, 1)	0.01 / Inf	1.87 [1.55, 2.24]
μ^{α_s}	N(10, 5)	1 / Inf	7.54 [6.89, 8.2]
τ^{α_s}	N(0, 1)	0.01 / Inf	2.36 [1.96, 2.82]
$\mu^{\beta_{confirmation}}$	N(0, 0.2)	-0.5 / 0.5	0.19 [0.13, 0.27]
$\tau^{\beta_{confirmation}}$	N(0, 0.2)	0.01 / Inf	0.23 [0.16, 0.29]
$\mu^{\beta_{proximity}}$	N(0, 0.2)	-0.5 / 0.5	0.02 [0, 0.04]
$\tau^{\beta_{proximity}}$	N(0, 0.2)	0.01 / Inf	0.05 [0, 0.08]
$\alpha^{\alpha_{keep}}$	N(1, 0.4)	0.01 / Inf	0.3 [0.22, 0.46]
$\beta^{\alpha_{keep}}$	N(5, 2)	0.01 / Inf	0.97 [0.66, 1.55]
$\mu^{\beta_{keep}}$	N(0, 0.5)	-Inf / Inf	-0.37 [-0.55, -0.26]
$\tau^{\beta_{keep}}$	N(0, 1)	0.01 / Inf	0.15 [0.05, 0.28]
$\alpha^{\alpha_{adopt}}$	N(1, 0.4)	0.01 / Inf	
$\beta^{\alpha_{adopt}}$	N(5, 2)	0.01 / Inf	
$\mu^{\beta_{adopt}}$	N(0, 0.5)	-Inf / Inf	
$\tau^{\beta_{adopt}}$	N(0, 1)	0.01 / Inf	

Table B5. Pearson correlations for parameter estimates across individuals and their 95% confidence intervals. Overall, we did not find strong correlations of parameter estimates, with three exceptions: (1) Individuals who put less weight on themselves (higher σ_p^2) tend to apply stronger proximity weighting (higher $\beta_{proximity}$) and (2) tend to have a lower tendency to keep their initial beliefs (lower α_{keep}). (3) The keep heuristic of individuals with a stronger tendency to keep their initial beliefs (higher α_{keep}), is less sensitive to the distance of the closest peer (less negative β_{keep})

	σ_p^2	α_s	$\beta_{confirmation}$	$\beta_{proximity}$	α_{keep}	β_{keep}
σ_p^2	1					
α_s	-0.02 [-0.22, 0.18]	1				
$\beta_{confirmation}$	-0.11 [-0.31, 0.09]	0 [-0.2, 0.2]	1			
$\beta_{proximity}$	0.35 [0.16, 0.51]	-0.15 [-0.34, 0.05]	-0.1 [-0.3, 0.1]	1		
α_{keep}	-0.22 [-0.41, -0.02]	0.03 [-0.17, 0.23]	0.11 [-0.1, 0.3]	0.04 [-0.16, 0.24]	1	
β_{keep}	-0.08 [-0.28, 0.12]	0.01 [-0.19, 0.21]	0.02 [-0.18, 0.22]	0.02 [-0.18, 0.23]	0.44 [0.26, 0.59]	1

Table B6. Description of the model parameters and the parameters of the their parent distribution (i.e., group-level priors).

Model feature	Parameter	Group-level prior	Description
Compromise			
Uncertainty own estimate	σ_p^2	$N(\mu^{\sigma_p^2}, \tau^{\sigma_p^2})$	Uncertainty associated with your own estimate
Intercept uncertainty peer estimate	α_s	$N(\mu^{\alpha_s}, \tau^{\alpha_s})$	The intercept uncertainty associated with the peer estimates
Distance weighting	$\beta_{confirmation}$	$N(\mu^{\beta_{confirmation}}, \tau^{\beta_{confirmation}})$	The influence of closeness of the peer estimate to the own estimate on the uncertainty associated with the peer estimate
Proximity weighting	$\beta_{proximity}$	$N(\mu^{\beta_{proximity}}, \tau^{\beta_{proximity}})$	The influence of proximity of the peer estimate to other peers on the uncertainty associated with the peer estimate
Keep			
Keep intercept	α_{keep}	$Beta(\alpha^{\alpha_{keep}}, \beta^{\alpha_{keep}})$	A value between zero and one describing the baseline probability of keeping the initial estimate.
Keep slope	β_{keep}	$N(\mu^{\beta_{keep}}, \tau^{\beta_{keep}})$	The influence of the distance of the closest peer on the probability to keep the initial estimate.
Adopt			
Adopt intercept	α_{adopt}	$Beta(\alpha^{\alpha_{adopt}}, \beta^{\alpha_{adopt}})$	A value between zero and one describing the baseline probability of adopting the estimate of a peer.
Adopt slope	β_{adopt}	$N(\mu^{\beta_{adopt}}, \tau^{\beta_{adopt}})$	The influence of distance on the probability to adopt the estimate of a peer.

Supplementary Methods

Definition of experimental conditions. Participants faced four experimental conditions in which they could adjust their initial estimates based on three pieces of social information. These conditions varied in the variance and skewness of this social information (Fig. 3.1E, main text). For each round, for each possible first estimate (E_1) we considered each possible triple of unique prerecorded estimates, and calculated the first three moments of its distribution (mean μ , variance σ^2 and skewness γ). To determine which triple would be shown in a given condition in a given round for a given value of E_1 , we used a cost function that assigned penalties (L) to deviations from the target mean (T), target variance (T_{σ^2}) and skewness (γ). For each round, for each possible value of E_1 we selected the triple with the lowest L. The cost functions L for each condition are given in the below table, which shows the penalties for deviations from the target mean, variance and skewness in separate columns:

Mean		Variance		Skewness
Low variance, no skew (LN)				
$100 \cdot \mu - T_\mu $	+	$10 \cdot \sigma^2 - T_{\sigma^2} $	+	$10 \cdot \gamma $
High variance, no skew (HN)				
$100 \cdot \mu - T_\mu $	+	$10 \cdot \sigma^2 - T_{\sigma^2} $	+	$500 \cdot \gamma $
High variance, leaning away from E_1 (HA)				
$100 \cdot \mu - T_\mu $	+	$10 \cdot \sigma^2 - T_{\sigma^2} $	-	$1,000 \cdot \gamma'$
High variance, leaning towards E_1 (HT)				
$100 \cdot \mu - T_\mu $	+	$10 \cdot \sigma^2 - T_{\sigma^2} $	+	$1,000 \cdot \gamma'$

In all conditions, T_μ was set to deviate 20% from E_1 . We held this distance fixed to avoid possible effects of the mean deviation of social information on its impact on behavior (Moussaid et al., 2013; Jayles et al., 2017; see also Fig. B6). For further standardization, T_μ was always in the direction of the true value A (Yaniv 2000; Molleman et al., 2019). Formally, $T_\mu = 1.2 \times E_1$ if ($E_1 > A$) and $T_\mu = 0.8 \times E_1$ if ($E_1 < A$). We set $T_{\sigma^2} = 10$ for the LN condition and $T_{\sigma^2} = 100$ for all other experimental conditions. For the LN and HN we aimed for symmetric distributions so we penalized positive absolute values of γ . For HA and HT, target skewness depended on whether E_1 was higher or lower than A. To this end, we used γ' which was equal to γ if $E_1 > A$ and equal to

$-\gamma$ if $E_1 < A$. This procedure ensured that participants faced well-defined experimental conditions (Fig. B7). For the 10 ‘filler’ rounds, which were intermixed with the experimental rounds, we randomly drew social information from the prerecorded pool.

We further implemented two control conditions completed in separate blocks of the experimental session (these blocks were completed in randomized order). First, participants completed trials in which they did not observe the stimulus themselves, but only observed the estimates of four peers (Fig. B2). The distribution of these peer estimates emulated the distributions of social information in each of the experimental conditions, enabling us to compare how individuals integrate personal and social information with a control in which individuals integrate four pieces of information, none of which is their own estimate. Second, participants could observe the estimate of only one peer whose deviation from the individuals’ first estimate matched that of the mean deviation in the four experimental conditions. The results from this one-peer control condition are not the focus of this paper and will not be reported here.

Cognitive model. To analyse potential strategies individuals use we developed a mixture model approach (Rieskamp et al., 2003), where individuals either apply ‘compromising’, ‘keeping’ or ‘adopting’.

Heuristic adjustment strategies of keeping and adopting. To account for ‘keep’ and ‘adopt’ assume that these two strategies are chosen with the mixture probabilities $P(keep)$ and $P(adopt)$ and compromising is chosen with probability $1 - (P(keep) + P(adopt))$. The keep probability is defined by a standard logistic function:

$$P(keep) = S(\text{logit}(\alpha_{keep}) + \beta_{keep} \times d), \tag{1}$$

where $\text{logit}(\alpha_{keep})$ and β_{keep} are the intercept and slope of the sigmoidal function, and d is the distance between the closest peer estimate and the first estimate ($d = |E_1 - X_1|$). Note that α_{keep} is a number between zero and one and transformed to a continuous scale via $\text{logit}(\alpha_{keep})$. This was done to account for individual differences with a Beta parent distribution (see Table B6). Similarly, the probability to ‘adopt’ the estimate of nearest neighbour depending on their distance is given by:

$$P(adopt) = S(\text{logit}(\alpha_{adopt}) + \beta_{adopt} \times d), \tag{2}$$

Compromising. We assume that each individual tries to judge the number of animals (N) for each image. The individual has two information sources: (i) The individuals initial belief about the number of animals (E_p) and social information provided by the peers (x_s) with $s = 1 : S$ and S being the group size. We assume the initial belief to have a discretized Normal probability distribution centered around first estimate (E_1) and associated with uncertainty (σ_p^2):

$$p(E_p|N) \sim \text{Norm}(E_1, \sigma_p^2). \quad (3)$$

Similarly, we assume the social information (SI_s) in the form of the provided peer estimates (X_s) and their associated uncertainty (σ_s^2) to be Normal distributed:

$$p(SI_s|N) \sim \text{Norm}(X_s, \sigma_s^2). \quad (4)$$

With only one peer (i.e. $S = 1$) we obtain the (posterior) probability of N animals applying Bayes rule:

$$p(N|E_p, SI_1) = \frac{p(E_p|N) \times p(SI_1|N)}{p(E_p, SI_1)} \quad (5)$$

Whereby $p(N|E_p, SI_1)$ is the new ‘posterior’ belief N (the number of animals) given own and social information (i.e., E_p and SI_1). Accordingly, with three peers the updating procedure is conducted with all peer estimates:

$$p(N|E_p, SI_1, SI_2, SI_3) = \frac{p(E_p|N) \times p(SI_1|N) \times p(SI_2|N) \times p(SI_3|N)}{p(E_p, SI_1, SI_2, SI_3)} \quad (6)$$

Note that the order is not affecting the generated probabilities. We distinguish between two features of how individuals weight social information or, more specifically, infer the uncertainty associated with a peer estimate (σ_s^2). First, individuals weight the peer estimates SI_s depending on their absolute distance to the individuals first estimate (d_s):

$$d_s = |E_1 - X_s|, \quad (7)$$

$$\sigma_s^2 = \alpha_s + \beta_{confirmation} \times d_s, \quad (8)$$

where α_s and $\beta_{confirmation}$ describe the intercept and slope, in other words, the overall sensitivity to social information and the influence of the distance, respectively. Similarly, individuals can weight peer estimates depending on their sum distance to the other peers (τ_s):

$$\sigma_s^2 = \alpha_s + \beta_{proximity} \times \tau_s, \quad (9)$$

or both:

$$\sigma_s^2 = \alpha_s + \beta_{confirmation} \times d_s + \beta_{proximity} \times \tau_s, \quad (10)$$

Note that for the fitting process, we centered the predictors d_s and τ_s their mean. Further, we set the minimal probability in the probability density function to be 10^{-30} to avoid outcome probabilities of zero.

Model fitting. We fitted the model using a hierarchical Bayesian inference technique implemented with “RStan” in R (R Core Team, 2019; Stan Development Team, 2018). We used a hierarchical structure where each parameter of the model has a higher order group-level prior (see Table B6). As default, these priors were normal distributions with hyper-parameters describing the mean and variance. The intercept parameters of the keep and adopt strategy are restricted between zero and one and are therefore described by a beta distribution with hyper-parameters and controlling the shape of the distribution. We ran 4 chains in parallel with 1,000 iterations each and discarded the first 500 as burn-in. We reduced the memory load by thinning the chains with a factor of 5. We investigated the importance of the four model features: (i) weighting depending on the distance of peer, (ii) weighting depending the peer on the proximity to other peers, (iii) keep heuristic and (iv) the adopt heuristic by calculating the leave-one-out cross-validation (looic; Vehtari et al., 2019) of the models compound of all possible combinations of these features (16 in total; see Table B3). We quantified the importance of a feature by calculating the average reduction of the looic when the feature was included (Fig. 3.2B). We report the fittings of the model with the lowest looic. Visual inspection of Markov chains and the Gelman Rubin statistic \hat{R} indicated that all Markov chains of all investigated 16 models converged.

Experimental materials

Participants were recruited from the crowdsourcing platform Amazon Mechanical Turk (MTurk). On that platform, our experiment was advertised as a ‘Human Intelligence Task’ (HIT) with a link that led to the experimental screens. Upon completion of the experiment, participants received a code that they could fill out on MTurk to receive their participation fee of \$4.50, plus a performance-specific bonus.

The experimental session consisted of 3 blocks (referred to ask ‘Task I, II and III’ in the instructions for participants), the order of which was randomized across participants.

Block 1: Participants observed an image and made their first estimate of the number of animals on it. Then, they could observe the estimate of three participants who completed the task before, and then make a second estimate. Participants completed five rounds for each of the experimental

conditions (Fig. 3.1E), plus 10 ‘filler rounds’ in which they observed three randomly drawn previous participants. So, in this block, participants completed 30 rounds in total; the participants’ responses in the experimental rounds of this block are the main focus of this study

Block 2: Like in Block 1, participants observed an image and made their first estimate of the number of animals on it, Then, they could observe the estimate of one participant who completed the task before, and then make a second estimate. This was repeated for five rounds (with a new image showing another species of animal in every round). NB: the data for this control condition is not the focus of this paper, and therefore its results are not presented here.

Block 3: Participants did not observe an image, but had to make an estimate based on the estimates of four participants who completed the task before. The five rounds of this task mimicked the four experimental conditions plus a random round (see Methods).

References

- Jayles, B., Kim, H.-r., Escobedo, R., Cezera, S., Blanchet, A., Kameda, T., Sire, C., and Theraulaz, G. (2017). How social information can improve estimation accuracy in human groups. *Proceedings of the National Academy of Sciences*, 114(47):12620–12625.
- Molleman, L., Kurvers, R. H., and van den Bos, W. (2019). Unleashing the beast: a brief measure of human social information use. *Evolution and Human Behavior*, 40(5):492 – 499.
- Moussaïd, M., Kämmer, J. E., Analytis, P. P., and Neth, H. (2013). Social Influence and the Collective Dynamics of Opinion Formation. *PLoS ONE*, 8(11):e78433.
- Nakagawa, S. and Schielzeth, H. (2010). Repeatability for gaussian and non-gaussian data: a practical guide for biologists. *Biological Reviews*, 85(4):935–956.
- R Core Team (2019). *R: A language and environment for statistical computing*. R Foundation for Statistical Computing, Vienna, Austria.
- Rieskamp, J., Busemeyer, J. R., and Laine, T. (2003). How Do People Learn to Allocate Resources? Comparing Two Learning Theories. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 29(6):1066–1081.
- Stan Development Team (2018). RStan: the R interface to Stan. R package version 2.18.2.
- Stoffel, M. A., Nakagawa, S., and Schielzeth, H. (2017). rptR: repeatability estimation and variance decomposition by generalized linear mixed-effects models. *Methods in Ecology and Evolution*, 8(11):1639–1644.
- Vehtari, A., Gabry, J., Yao, Y., and Gelman, A. (2019). loo: Efficient leave-one-out cross-validation and waic for bayesian models. R package version 2.1.0.

C | Supplementary Material to Chapter 4

Supplementary Results

Using the Two-stage Dynamic Signal Detection (2DSD) model to model the personal phase: To model the personal choice phase, we used the two-stage dynamic signal detection (2DSD) model. The 2DSD model is an evidence accumulation model that can account for choice and response time (RT) in the personal choice and the associated confidence judgement. In so doing, it can identify cognitive mechanisms potentially governing the interrelationships of these behavioral measures (Pleskac and Busemeyer, 2010). Like other evidence accumulation models, it assumes that individuals gather evidence over time until the amount of evidence surpasses a threshold. The two key assumptions of the 2DSD model are that evidence accumulation continues after the decision is made and that reported confidence depends on the evidence accumulated at the time of the confidence judgement. Thereby, the evidence state is mapped into confidence judgements using response criteria that serve as thresholds indicating the next higher confidence judgements (e.g., from 50 to 60, or 60 to 70). See Pleskac & Busemeyer (2010) for a detailed description of the 2DSD model.

We fitted the model in the hierarchical Bayesian framework, implemented with RStan in R (R Core Team, 2019; Stan Development Team, 2018), with five parallel chains with 10,000 iterations each and a thinning factor of 10. The first half of the iterations were discarded as burn-in. Descriptions of the main parameters are given in Supplementary Table C6. For the Wiener diffusion process, we included boundary separation α , predecisional drift rate δ_{pre} , relative start point z , and nondecision time NDT , which was calculated relative to the fastest response. Some trials were expected to be more difficult than others, because the number of sharks could be closer to (i.e., 4 and 6) or further away from (i.e., 3 and 7) the threshold number of sharks (5). We accounted for variations in difficulty by varying the predecisional drift rate δ_{pre} , depending on trial difficulty:

$$\delta_{pre} = \begin{cases} \delta_{difficult}, & \text{if 4 or 6 sharks present} \\ \delta_{difficult} + \Delta_{easy}, & \text{if 3 or 7 sharks present} \end{cases} \quad (1)$$

with Δ_{easy} describing the additional effect of easy trials on the drift rate. For the postchoice process, we fitted confidence criteria and the postdecisional drift rate (δ_{post}). The postdecisional drift rate is influenced by the predecisional drift rate, with the parameter w controlling its strength, and δ_{choice} describing the influence of the choice on the subsequent drift:

$$\delta_{post} = \begin{cases} w \times \delta_{pre} + \delta_{choice}, & \text{if correct} \\ w \times \delta_{pre} - \delta_{choice}, & \text{if incorrect} \end{cases} \quad (2)$$

The evidence distribution at the time point when confidence is reported L_{conf} is a combination of the evidence accumulated at the time point of choice and the evidence accumulated between choice and confidence judgement. It is normally distributed with a mean of

$$E[L_{conf}] = \begin{cases} \alpha + \delta_{post} \times IJT, & \text{if correct} \\ 0 + \delta_{post} \times IJT, & \text{if incorrect} \end{cases} \quad (3)$$

and a variance of

$$var[L_{conf}] = \sigma^2 IJT \quad (4)$$

with IJT being the interjudgement time (i.e., the time between choice and confidence reporting).

Each decision maker has confidence criteria c_j to map the evidence state L_c into six possible confidence judgements $conf_j$ with $j = 0, 1, 2, \dots, 5$, corresponding to the confidence levels 50 to 100. The probability of reporting $conf_j$ is given by the normal cumulative distribution $\sim N(E[L_c], var[L_c])$ with:

$$P(c_j < L_c < c_{j+1}) \quad (5)$$

where c_0 is equal to $-\infty$ and c_6 to ∞ . The five remaining confidence criteria are fitted by the model. We assume the locations of the confidence criteria for correct and incorrect choices to be symmetrical. Hence, we set the locations relative to the choice thresholds with $alpha + c_j$ and $0 - c_j$ for correct and incorrect choices, respectively. For the fitting process, we excluded all choices with RTs below 0.1 sec. To compare the predictions of the model with the empirical data, we generated choices, RTs, and confidence judgements using the participant’s mean posterior parameter estimate. The confidence judgements were generated by sampling from the evidence distribution at the time point of the judgement and mapping this evidence state to a confidence judgement.

We thus obtained confidence judgements given the individual’s choice, RT, and interjudgement time. To account for stochasticity generated by the sampling process, we sampled 100 confidence judgements, choices, and RTs per individual and trial.

2DSD model results: Participants drifted towards the correct choice threshold ($\delta_{difficult} = 0.37$, CI = [0.33, 0.41]). Trials with three or seven sharks were easier than trials with four or six sharks, as indicated by a stronger drift towards the correct option in the former ($\Delta_{easy} = 0.05$, CI = [+0.00, 0.10]). Varying drift rates depending on difficulty were not included in the social DDM analysis, as the effect was comparatively small. After making a choice, participants continued accumulating evidence and, on average, kept gathering correct evidence ($w = 0.72$, CI = [0.62, 0.83]). Hence, participants who made an incorrect decision gathered more evidence over time contradicting their initial choice (resulting in lower confidence), whereas the evidence of those who made a correct choice was strengthened (resulting in higher confidences). This process explains the increasing difference in confidence ratings between correct and incorrect choices as interjudgement time increases (Supplementary Fig. C4a). Additionally, there was a choice effect on the postdecision drift, whereby participants accumulated evidence in favour of their already chosen option ($\delta_{choice} = 1.64$, CI = [1.47, 1.80]). As a result, longer interjudgement times are predicted to lead to higher confidence judgements (Supplementary Fig. C4b). Figure C4c shows that the 2DSD recreates the well-established relationship between confidence and accuracy, which is partly determined by the postdecisional processing evident in Figures C4a and b. In both the 2DSD and the social DDM analysis, we thus found that confidence is linked to the evidence state and that participants drifted in the direction of their chosen option (i.e., reinforced their ‘belief’ in their original choice). Figure C4d shows RT distributions for the personal choice. Overall, the empirical data (solid lines) correspond closely with the predictions of the 2DSD model (dashed lines), indicating that the personal phase can be described by a drift diffusion process. One distribution characteristic the model cannot recover is the higher average RTs for incorrect choices. This is a well-known property of the drift–diffusion model, and can be addressed by adding trial-by-trial variability to the drift rates (Ratcliff and Rouder, 1998). For simplicity, we have not included trial-by-trial variability.

Supplementary Figures




Figure C1. Improvement during the social phase depended on the quality of social information. Participants' choices were increasingly likely to improve/worsen as the size of the majority for the correct/incorrect option increased. Dots represent the mean; error bars represent twice the standard error. The dashed line shows the prediction of the social DDM.

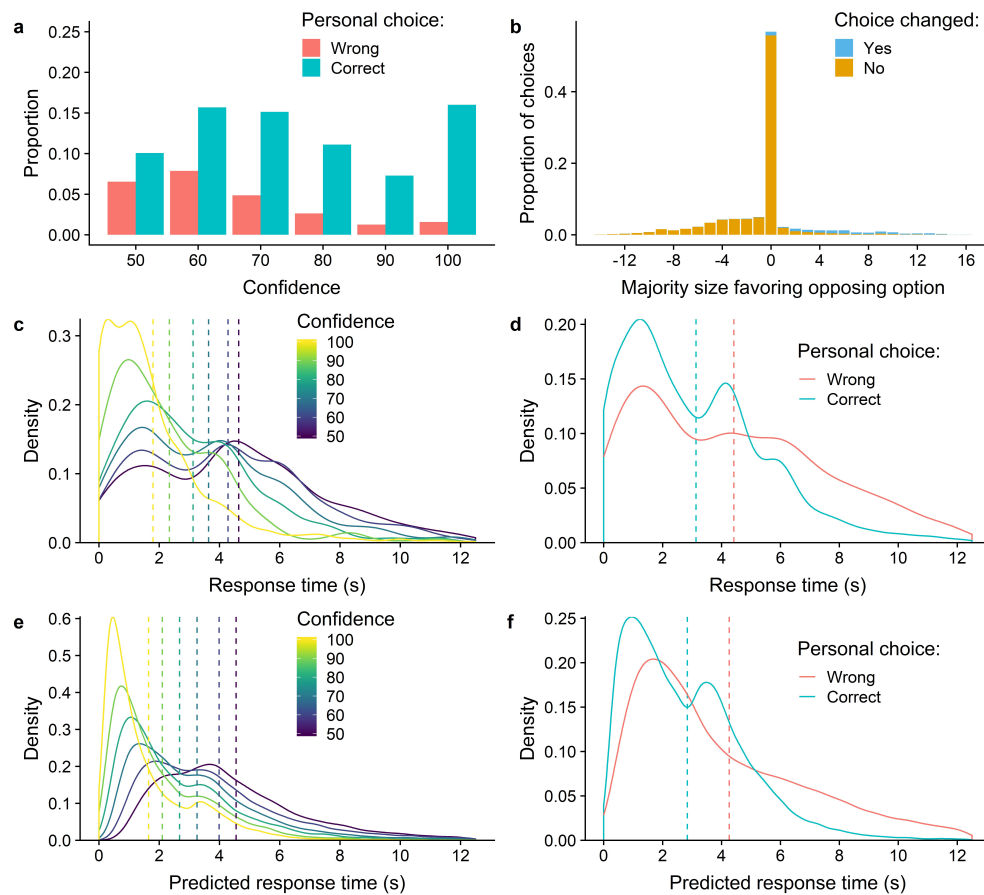


Figure C2. Distributions of key behavioural measures. (a) The proportion of reported confidence scores for correct and incorrect choices. The higher the confidence score, the larger the proportion of correct choices, resulting in a positive confidence–accuracy relationship (see also Supplementary Fig. C4c). (b) The proportion of choices made in the presence of different majority sizes. In the social phase, most choices ($\approx 60\%$) were made in the absence of a majority, and participants who experienced a majority were more likely to observe a confirming majority (i.e., negative values) than an opposing majority. Participants facing an opposing majority were more likely to change their choice the larger the size of this opposing majority. (c) Observed RT distributions during the social phase as a function of reported confidence. Participants reporting the highest level of confidence overwhelmingly responded within 4 seconds, whereas the distribution of participants reporting the lowest confidence level peaked after 4 seconds. (d) Observed RT distributions during the social phase for correct and incorrect choices. Given that unconfident participants are more likely to be wrong and waited longer, it follows that individuals who were wrong, on average, wait longer to respond. (e, f) RT distributions as predicted by the social DDM. The model recovers not only the relationship of RT with confidence and accuracy, but also the shape of the RT distributions. The RT distributions are multimodal because social information was first updated after 3 seconds and then every 2 seconds. The updating events often resulted in larger majorities which increased the likelihood of a response by an increase in the drift rate. (c–f) Dashed vertical lines represent the mean RTs.

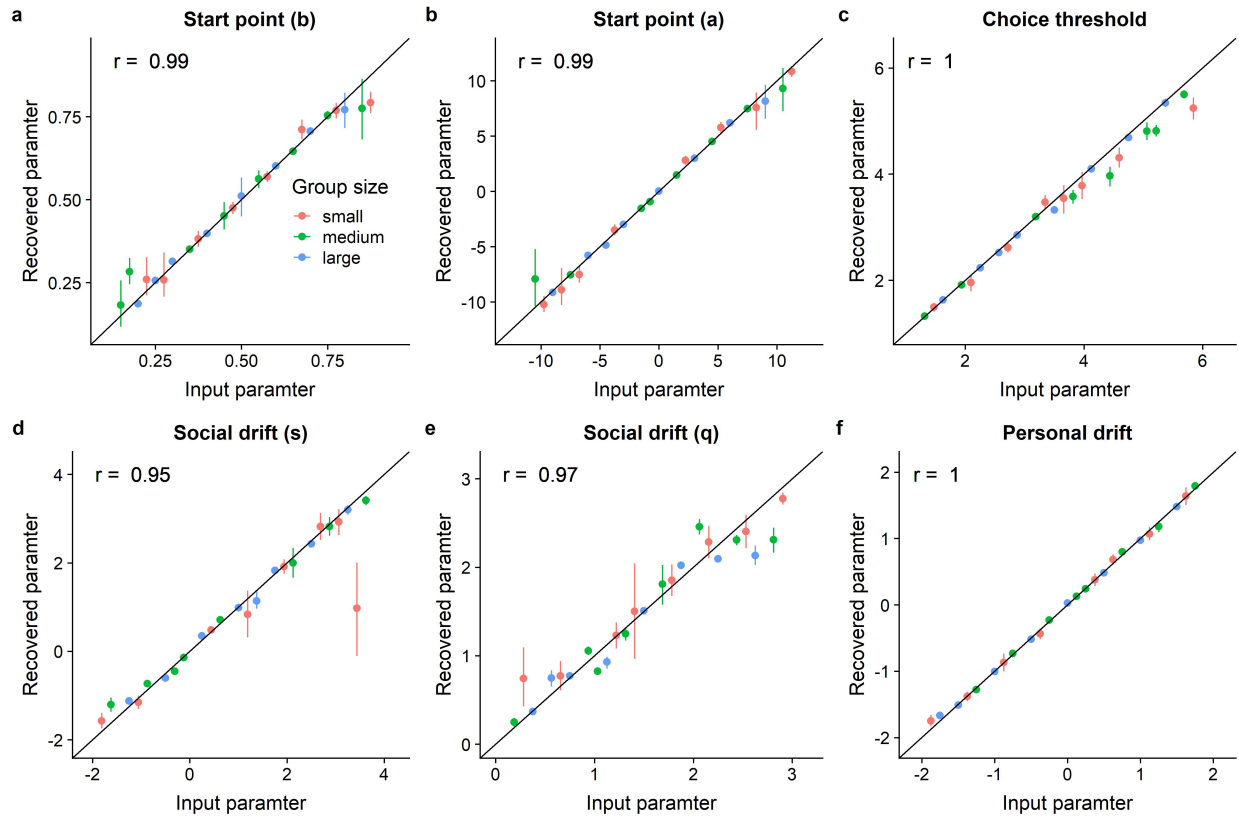


Figure C3. Model recovery. The x-axis shows the actual (input) parameters; the y-axis shows the recovered parameters. The figure shows the results of a parameter recovery analysis conducted to ensure that the parameters of the social DDM are interpretable and capture distinct cognitive mechanisms. We repeatedly generated data with random input parameters and recovered them with the same hierarchical social DDM used to analyse the empirical data. The input parameters were sampled with a quasirandom number generator (using the sobol sequence), ensuring an even distribution across a large multidimensional parameter space. Using these input parameters, we generated social choices by computing probability density functions while taking into account the personal choice, reported confidence, and the social information observed by the participant at a given trial. The generated data thus have the same hierarchical structure as the empirical data, with 141 participants and varying group size. Again, we report the mean of the posterior distributions and the 95% CI of the higher order group-level estimate for each group size. To measure the relationship of input and recovered parameters, we calculated Spearman’s correlation coefficient r for all parameters (except nondecision time, which is relative to a participant’s fastest response and thus meaningless on a group level). For all parameters, there was a strong positive correlation between the generated and the recovered parameters. The estimates provided by the social DDM thus describe separate identifiable features and are interpretable in their magnitude.

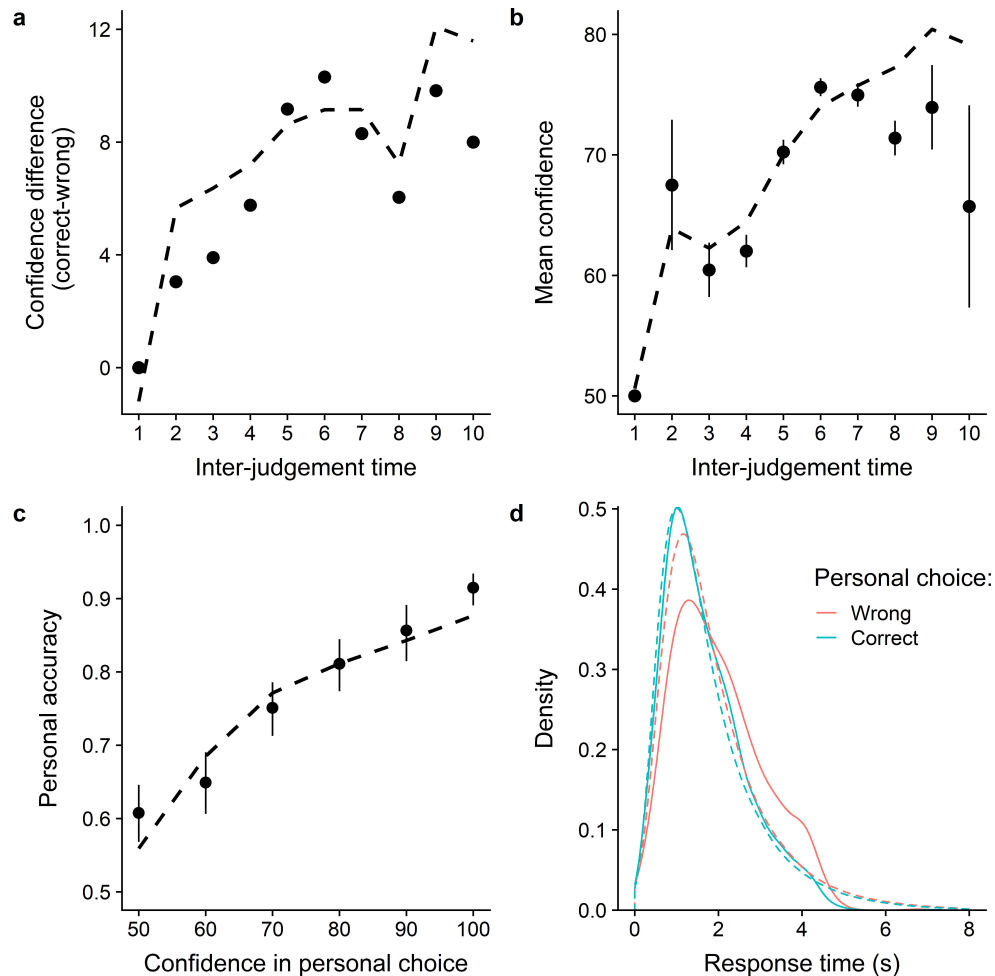


Figure C4. Results of the 2DSD model. (a) The longer the time between the personal choice and the confidence judgement (interjudgement time), the larger the difference in confidence between participants whose choices were correct vs. incorrect. Dots represent the average confidence judgements for correct choices minus the average confidence judgements for incorrect choices for different interjudgement times. (b) The longer the interjudgement time, the higher the reported confidence judgements. Dots represent the mean; error bars represent twice the standard error. (a–b) For visualization purposes, interjudgement times are binned by rounding to the closest integer. (c) Participants reporting higher confidence were more likely to be correct. Dots and error bars represent mean and 95% CI of the posterior distribution. (d) The solid lines represent the observed RT distribution of the personal choice for correct (blue) and incorrect (red) choices. (a–d) The dashed lines represent the predictions of the 2DSD model.

Supplementary Tables

Table C1. Bayesian linear regression results

Response Predictor	Estimate	Est.Error	l-95% CI	u-95% CI	Eff.Sample	Rhat
Accuracy						
Intercept (personal choice)	1.1	0.07	0.97	1.23	9657.34	1
Social choice	0.3	0.05	0.2	0.39	32162.95	1
Accuracy						
Intercept	-1.58	0.17	-1.91	-1.25	20461.1	1
Confidence	3.82	0.24	3.35	4.28	20270.94	1
Accuracy						
Intercept (personal choice)	1.65	0.08	1.48	1.81	7563.13	1
RT	-0.16	0.01	-0.18	-0.14	15170.53	1
RT: social choice	0.11	0.01	0.09	0.13	18797	1
Likelihood to change						
Intercept	-3.6	0.18	-3.96	-3.26	7735.1	1
Size of opposing majority	0.62	0.03	0.57	0.67	12889.8	1
RT						
Intercept	6.96	0.23	6.49	7.41	4644.56	1
Confidence	-4.86	0.18	-5.22	-4.5	9740.24	1
Improvement						
Intercept	1.17	0.2	0.78	1.56	22984.38	1
Confidence	-4.27	0.31	-4.88	-3.68	21332.94	1
Improvement						
Intercept (earlier; more accurate)	0.09	0.01	0.08	0.11	3830.29	1
Earlier; less accurate	-0.09	0	-0.09	-0.08	19689.95	1
Later; more accurate	-0.04	0.01	-0.05	-0.03	16925.55	1
Later; less accurate	-0.04	0.01	-0.06	-0.02	17547.86	1

Table C2. Deviance information criterions (DIC) for different versions of the social DDM. The version with the lowest DIC is indicated in bold.

	No further drift	Drift towards initial choice	Drift towards correct
Neither	79493	76026	77854
Varying start point	76364	74183	74851
Social drift	78058	74275	77286
Both	74200	71865	73835

Table C3. Mean parameter estimates and 95% credible intervals of the social DDM for different group sizes.

	Small	Medium	Large
NDT (T_s)	0.4 [0.23, 0.56]	0.33 [0.25, 0.41]	0.31 [0.26, 0.37]
Start point (a)	4.2 [3.11, 5.35]	3.43 [2.81, 4.07]	3.9 [3.46, 4.37]
Start point (b)	0.5 [0.45, 0.54]	0.48 [0.45, 0.51]	0.5 [0.48, 0.52]
Personal drift	0.65 [0.45, 0.86]	0.62 [0.5, 0.75]	0.53 [0.47, 0.59]
Social drift (s)	0.51 [0.23, 0.82]	0.31 [0.24, 0.38]	0.36 [0.3, 0.41]
Social drift (q)	1.75 [1.16, 2.36]	0.93 [0.81, 1.05]	0.66 [0.6, 0.72]
Choice threshold	3.22 [2.58, 3.9]	3.43 [3.09, 3.77]	3.3 [3.04, 3.56]

Table C4. Differences between parameter estimates for different group sizes. Shown are the mean and the 95% credible intervals.

	Small – Medium	Small – Large	Medium – Large
NDT (T_s)	0.06 [-0.12, 0.25]	0.08 [-0.09, 0.26]	0.02 [-0.08, 0.11]
Start point (a)	0.77 [-0.48, 2.08]	0.3 [-0.88, 1.51]	-0.47 [-1.25, 0.31]
Start point (b)	0.01 [-0.04, 0.07]	-0.01 [-0.06, 0.04]	-0.02 [-0.06, 0.02]
Personal drift	0.03 [-0.2, 0.27]	0.12 [-0.08, 0.33]	0.09 [-0.04, 0.23]
Social drift (s)	0.2 [-0.09, 0.51]	0.15 [-0.13, 0.46]	-0.05 [-0.13, 0.04]
Social drift (q)	0.82 [0.22, 1.44]	1.1 [0.5, 1.71]	0.27 [0.14, 0.41]
Choice threshold	-0.21 [-0.93, 0.54]	-0.08 [-0.77, 0.65]	0.13 [-0.3, 0.56]

Table C5. The number of groups per group size.

Group size	Number of groups	Number of participants	Classification
3	5	15	Small
7	3	21	Medium
8	1	8	Medium
9	1	9	Medium
10	1	10	Medium
15	3	45	Large
16	1	16	Large
17	1	17	Large
Total:	16	141	

Table C6. Description of the parameters of the 2DSD.

Model feature	Parameter	Description
Nondecision time	NDT	A parameter between 0 and 1 accounting for nondecision time (e.g., motor response time), parameterized as the time relative to an individual’s fastest response.
Relative start point	z	Describes the initial evidence state before the evidence sampling process begins.
Predecisional drift rate	$\delta_{pre} = \begin{cases} \delta_{difficult}, & \text{if difficult} \\ \delta_{difficult} + \Delta_{easy}, & \text{if easy} \end{cases}$	The baseline predecisional drift rate for difficult (i.e., 4 or 6 sharks) and easy (i.e., 3 or 7 sharks) trials.
Boundary separation	α	The boundary separation determines how much evidence an individual has to accumulate to make a decision.
Carryover effect	w	A parameter controlling how strongly the predecisional drift rate carries over to the postdecisional drift rate.
Self-confirmation bias	δ_{choice}	A parameter describing the influence of the choice (i.e., being correct or incorrect) on the subsequent drift rate.
Confidence criteria	c_j	Thresholds that divide the evidence space into confidence judgements.

Table C7. 2DSD parameter results

Parameter	Estimate	l-95% CI	u-95% CI	Eff.Sample	Rhat
Nondecision time	0.63	0.56	0.74	2362.88	1
Relative start point	0.53	0.52	0.54	2085.7	1
Predecisional drift rate (intercept, difficult)	0.37	0.33	0.41	1991.81	1
Predecisional drift rate (effect of easy)	0.05	0	0.1	2256.05	1
Boundary separation	2.5	2.45	2.56	2268.78	1
Carryover effect	0.72	0.62	0.83	2354.22	1
Self-confirmation bias	1.64	1.47	1.8	1962.74	1
Confidence criteria 5	3.27	2.46	4.09	1875.05	1
Confidence criteria 4	4.99	4.54	5.44	2444.76	1
Confidence criteria 3	3.83	3.5	4.17	2160.56	1
Confidence criteria 2	3.04	2.74	3.35	2499.86	1
Confidence criteria 1	2.4	2.11	2.72	2605.37	1

References

- Pleskac, T. J. and Busemeyer, J. R. (2010). Two-stage dynamic signal detection: A theory of choice, decision time, and confidence. *Psychological Review*, 117(3):864–901.
- R Core Team (2019). *R: A language and environment for statistical computing*. R Foundation for Statistical Computing, Vienna, Austria.
- Ratcliff, R. and Rouder, J. N. (1998). Modeling response times for two-choice decisions. *Psychological Science*, 9(5):347–356.
- Stan Development Team (2018). RStan: the R interface to Stan. R package version 2.18.2.

D | Supplementary Material to Chapter 5

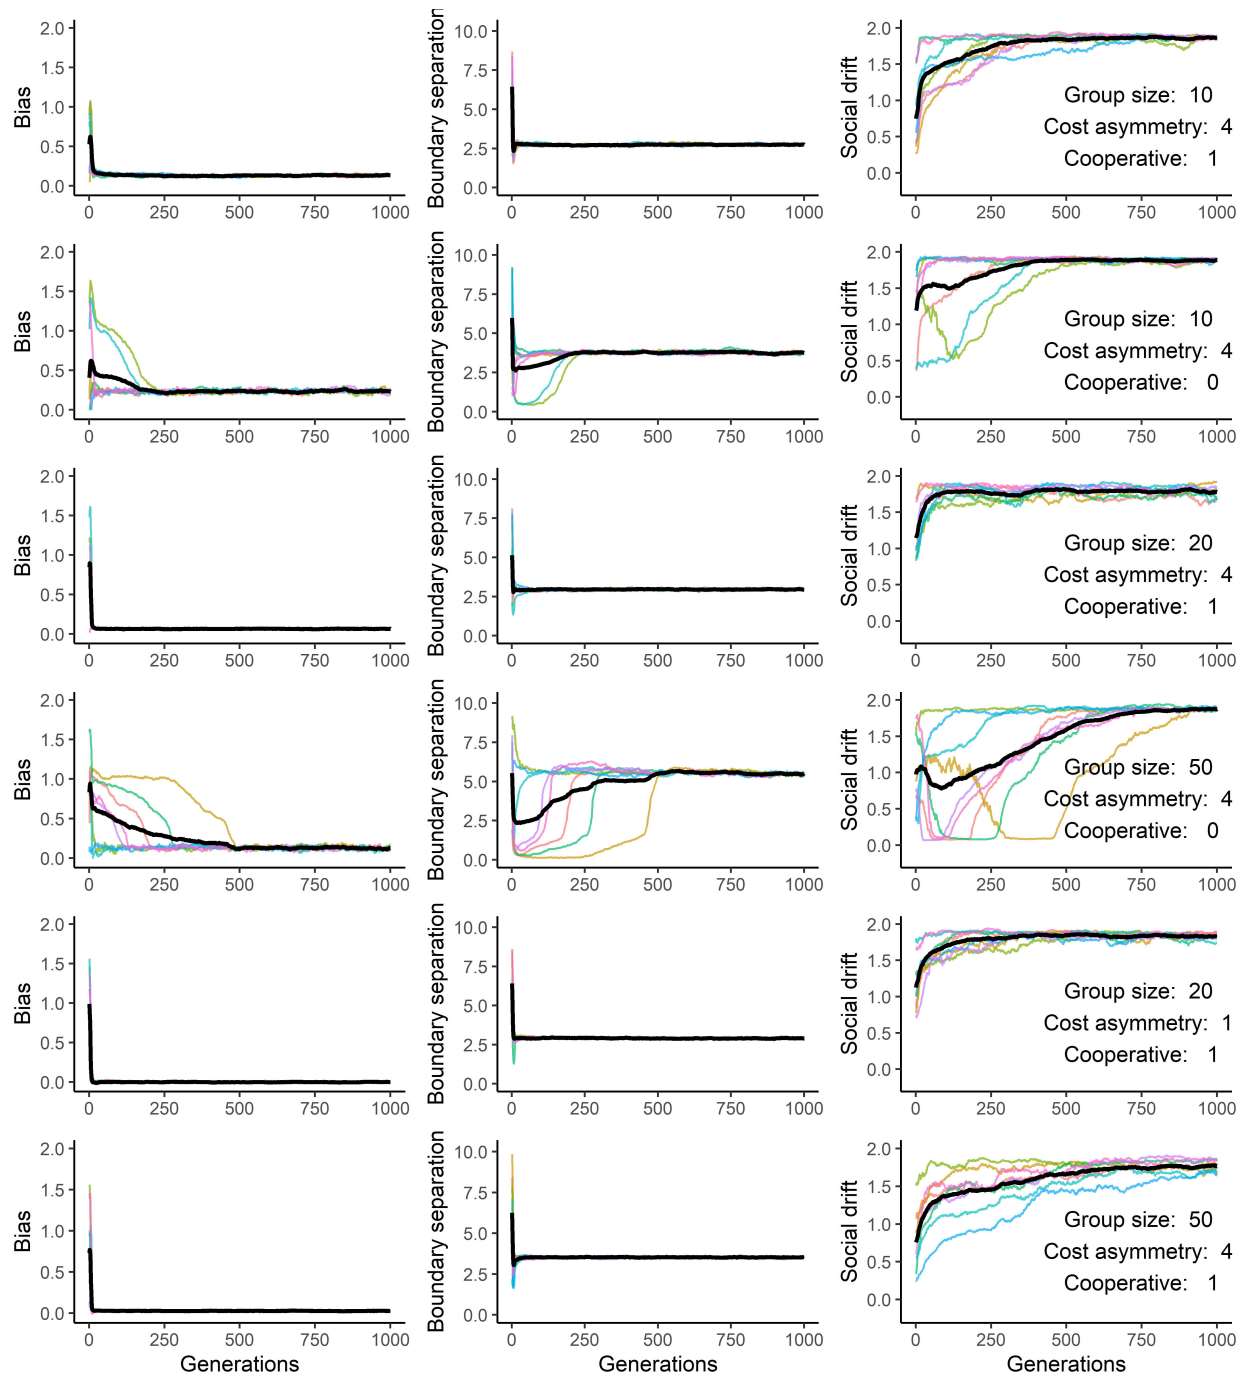


Figure D1. Example trajectories of the evolutionary algorithm. Shown are the evolutionary trajectories of bias (left), boundary separation (center), and social drift rate (right), for six additional scenarios. These scenarios were randomly drawn from all 30 analysed scenarios. The corresponding parameter settings are shown in the right panels. The colored lines represent the average parameter value within each of the eight evolving populations and the black line indicates the average across all eight populations.

Declaration of Independent Work

I hereby declare that:

- I completed this doctoral thesis independently. Except where otherwise stated, I confirm that the work presented in this thesis is my own.
- Where information has been derived from other sources, I confirm that this has been indicated in the thesis.
- I have not applied for a doctoral degree elsewhere and do not have a corresponding doctoral degree.
- I have acknowledged the Doctoral Degree Regulations which underlie the procedure of the Department of Education and Psychology of Freie Universität Berlin, as amended on August 8th 2016.
- The principles of Freie Universität Berlin for ensuring good academic practice have been complied with.

Alan Novaes Tump

Berlin, 18. December 2019