



rijksuniversiteit
 groningen

Social learning and human cooperation

PhD thesis

to obtain the degree of PhD at the
University of Groningen
on the authority of the
Rector Magnificus Prof. E. Sterken
and in accordance with
the decision by the College of Deans.

This thesis will be defended in public on
Friday 21 February 2014 at 16.15 hours

by

Lucas Sebastiaan Molleman

born on 28 August 1984
in Lichtenvoorde

Supervisor

Prof. F.J. Weissing

Assessment committee

Prof. D. Engelmann

Prof. A.G.G. Groothuis

Prof. M. Milinski

The research reported in this thesis was carried out at the Theoretical Biology Group, which is part of the Centre for Ecological and Evolutionary Studies of the University of Groningen (The Netherlands), according to the requirements of the Graduate School of Science (Faculty of Mathematics and Natural Sciences, University of Groningen).

Lucas Molleman received an Ubbo Emmius grant from the Faculty of Mathematics and Natural Sciences.

Part of this work is funded by the Institute of Biodiversity and Ecosystem Dynamics of the University of Amsterdam (The Netherlands) and the Dr. J.L. Dobberke Foundation for Comparative Psychology.

The printing of this thesis was partly funded by the University of Groningen and the Faculty of Mathematics and Natural Sciences.

Layout and figures: Lucas Molleman

Cover design: Lucas Molleman

Printed by: GVO Drukkers & Vormgevers | Ponsen & Looijen

ISBN: 978-90-367-6811-5 (printed version)

ISBN: 978-90-367-6845-0 (electronic version)

Index

Chapter one	
General introduction	5
Chapter two	
Effects of conformism on the cultural evolution of social behaviour	23
Chapter three	
Cultural evolution of cooperation: the interplay between forms of social learning and group selection.....	51
Chapter four	
Consistent individual differences in human social learning strategies	73
Chapter five	
Time consistency in human social learning strategies and the effects of their assortment on cooperative interactions.....	129
Chapter six	
Personal experience and reputation interact in human decisions to help reciprocally	149
Chapter seven	
Diversity of morals in human cooperation	179
Chapter eight	
General discussion	211
References	225
Dutch summary	239
Acknowledgements	251
Biography	253

Chapter one

General introduction

Humans largely owe their dominant role on Earth to their social nature. The cooperation and coordination of actions in groups has enabled humans to surmount adaptive challenges that would be impossible for any individual to cope with on its own. Whereas cooperative action appears to be an obvious part of human life, it is a puzzling phenomenon from an evolutionary viewpoint. Namely, natural selection generally favours selfish over cooperative behaviour. Nevertheless, in countless kinds of social contexts, human interactions lead to outcomes that benefit all parties involved. Both in human prehistory and in modern day, sophisticated forms of team work has facilitated fruitful outcomes of human collective action: from cooperative hunting of mammoths on the Pleistocene tundra, to the coordinated creation of high-tech products in a modern factory.

In addition, humans are champions of social learning. People pass on useful information through teaching and often improve their behaviour by observing others. The transmission of socially learned information can change the distribution of behaviours in a population. Compared to genetic evolution, this process of 'cultural evolution' through social learning can enable a population to rapidly adjust to the environment as adaptive information spreads between its members. Cultural evolution can also have a cumulative character: over time, populations can build up large bodies of adaptive knowledge that individuals can use and refine. Gradually, this process generated technologies that enhanced the opportunities to adapt to a vast array of habitats across the world, and culminated in the largely self-constructed environments we inhabit today.

Cultural evolution also shapes human social behaviour. In determining their behaviour in social contexts, people attend to the behaviour of their peers, and regard the social norms established in the population. The direction and outcome of cultural evolution depends on the forms of social learning that people use; depending on those forms of social learning, cultural evolution of social behaviour can lead to outcomes different from genetic evolution by natural selection.

In this thesis, I investigate the interplay between human cooperation and social learning. First, I develop a set of conceptual models to map out how various forms of social learning affect the direction and outcome of

cultural evolution in a range of interaction contexts (Chapters 2 and 3). Second, with a series of decision making experiments, I assess what forms of social learning people use when they are confronted with such contexts (Chapters 4 and 5). Third and finally, I turn to the question of how people condition their cooperative behaviour on past actions of their interaction partners, shedding light on how human cooperation can be sustained by reciprocal helping (Chapters 6 and 7).

The puzzle of cooperation

In numerous forms of social interaction, humans cooperate in groups to the benefit of all group members. No other animal can successfully coordinate efforts in such a vast array of social situations, ranging from helping to raise each other's kids to hunting big game, and from tending for the sick to defending a village against outside threats (*e.g.*, (Boyd and Richerson 1985; Mace and Sear 2005; Stiner et al. 2009). The benefits produced by such forms of cooperation broadened the range of environments in which humans can survive, thereby allowing our species to colonise virtually all terrestrial habitats around the world. Over evolutionary time, humans did not only vastly grow in numbers, but also societies have grown from hunter-gatherer societies consisting of a few tens of people, to large nation states harbouring millions of inhabitants living together without being continuously in conflict with each other.

Despite the variety of different contexts in which humans cooperate, scientists across disciplines tend to focus their studies of social interaction on situations in which the interests of an individual are opposed to the interests of the group. In these so-called 'social dilemmas', cooperative behaviour comes with a cost for the actor and benefits others. Evolutionary theory, based on natural selection through differential reproduction, predicts that in such situations cooperation is not a straightforward outcome. Take, for example, the defence of a tribal village against a threat from the outside. The men who step forward and risk their lives to protect their tribe have a higher chance of being injured or even killed due to their bravery. All other things being equal, these men will have fewer chances to pass on their brave nature to future

generations, compared to the more cautious and cowardly members of their tribe. Evolutionary theory thus predicts that in the absence of other selective pressures, bravery tends to disappear over time due to differences in reproductive success (Darwin 1859; Darwin 1871). When social traits such as bravery or a tendency to cooperate are passed on by social learning rather than by genetic inheritance, the direction and outcome of evolution might be different. This subject will be discussed at length below. In this section, my main focus will be on natural selection.

Whereas the logic of natural selection seems to dictate that the selfish will prevail, costly cooperation is found not only in humans, but in many species throughout nature. For instance, when the bacterium *Pseudomonas aeruginosa* suffers from shortage in iron (which it needs to survive), it can secrete substances called 'siderophores' into the environment to bind otherwise insoluble iron, making the iron available for bacterial uptake. Producing these substances is costly for the individual bacterium, but all bacteria in the neighbourhood can take up the iron molecules when bound to siderophores (Guerinot 1994). Mutant bacteria that do not produce siderophores can reap the benefits of the cooperative efforts of others without paying the costs; in the light of natural selection, it is surprising that these siderophore-producing types are not driven to extinction. Another example of such costly cooperation is given by Belding's ground squirrels (*Spermophilus beldingi*), who exclaim alarm calls when a predator approaches their group (Sherman 1977). These calls alert others, but are risky for the caller who attracts the attention of the predator and runs the risk of being injured, or even killed. Individuals that refrain from emitting alarm calls can benefit from the alarm calls of others, and will have relatively better chances to survive and reproduce than their group mates that do take the risk. Emitting alarm calls thus decreases the fitness of an individual compared to its group mates, and is as such disfavoured by natural selection. In both these examples, costly cooperation (producing siderophores and emitting alarm calls) tends to lead to a lower reproductive success for an individual but benefits the group. This presents the puzzle of evolution of cooperation by means of natural selection: how can costly cooperation emerge and persist through evolutionary time when 'cheaters' can benefit from cooperation without paying the costs? In the following paragraphs, I will

briefly review a few factors that may explain the evolution of cooperation despite immediate selection against it.

Darwin (1871, chapter 5) already realised that cooperation poses fundamental problems for his theory, and he proposed a number of mechanisms that could explain how costly cooperation could be favoured over selfish behaviour. Currently, the mechanisms mentioned by Darwin are still at the core of research on the evolution of cooperation (see, for example, (Lehmann and Keller 2006) for an overview). Darwin pointed out that natural selection can take place not only between individuals, but also between groups (Darwin 1871), chapter 5). Throughout human evolutionary history, different tribes have constantly been in conflict with each other, and archaeological and ethnological evidence suggests that warfare contributed greatly to the total mortality of our ancestors in hunter-gatherer populations (Darwin 1871; Soltis et al. 1995; Bowles 2009). It is plausible that tribes with a lot of brave (*i.e.*, cooperative) individuals have been more likely to win conflicts and to supplant other tribes, or able to get control over the most fertile pieces of land that could sustain a larger number of people. The importance of group selection in human prehistory and its effects on the evolution of cooperation are, however, controversial issues. In Chapters 2 and 3 of this thesis I take up this subject, and discuss potential effects of group selection in the light of cultural evolution.

Darwin (1871) also recognised that reciprocity might lead to cooperation: when individuals interact repeatedly, a cooperative act might be returned at a later point in time. Reciprocal helping can lead to cooperation among pairs of individuals: scratching someone's back may pay off if he scratches yours at some later point in time (Trivers 1971). Human social interactions frequently involve forms of such 'direct reciprocity', and this mechanism is believed to occur in some animal species as well (Dugatkin and Reeve 1997; Clutton-Brock 2009). Reciprocal cooperation can also work in a more indirect manner. Darwin noted that humans are particularly sensitive to the praise and blame of their fellows (Darwin 1871), p. 164). Interactions between individuals are often observed by others, and people like to talk about the behaviour of their peers. Through eavesdropping and gossip, social information propagates through human groups such that individuals build up a reputation of being nice and

friendly, or nasty and selfish. When, as a consequence, individuals with a good reputation are more likely to receive help, indirect reciprocity can promote cooperation (Alexander 1987). Models suggesting that cooperation can be supported by reputation-based reciprocity, however, often reflect a rather simplistic view of how individuals account for past behaviour about their peers to allow for mathematical tractability (*e.g.*, (Nowak and Sigmund 1998; Panchanathan and Boyd 2003; Nowak and Sigmund 2005). Also, and the mechanisms of direct and indirect reciprocity are typically studied in isolation. In Chapters 6 and 7 of this thesis I will show experimentally how people react to historical information about the cooperativeness of their interaction partners (stemming from both direct experience and reputations), and assess how accounting for peers' motivations behind decisions making affects reputation-based cooperation.

Many decades after Darwin's books, researchers have proposed additional mechanisms that can help explain how natural selection can favour cooperation over selfishness. These mechanisms do not take the centre stage in this thesis, but it is worthwhile to highlight a some of them here since they play a role in discussions related to the topic. Partner choice is one of them (Noë and Hammerstein 1994): when individuals can choose who to interact with, cooperators can avoid defectors and exclude them from the benefits arising from their cooperative efforts. Interaction partners may even specialise in specific tasks and both profit from the exchange of products, increasing the benefits of mutual cooperation. Another mechanism that can promote cooperation is kin selection (Hamilton 1964; Maynard Smith 1964), in which cooperation can be favoured if its costs are compensated by the benefits for related recipients. It has been hypothesized that this mechanism works in *P. aeruginosa*, where higher levels of relatedness lead to higher levels of siderophore production (Griffin et al. 2004). Since cooperation is channelled mainly towards kin, the costs it entails to an individual are compensated for by enhanced benefits for its relatives. Another class of mechanisms that can promote cooperation involve punishment. When individuals can be punished if they fail to cooperate, the advantages of selfish behaviour can be offset, making cooperation favoured through force. Physical punishment and diverse forms of social exclusion are widely used in

human societies (Guala 2012; Van den Berg et al. 2012), partly to 'correct' those who fail to cooperate and harm the interests of society. Punishment of defectors also occurs in other species. For instance, in many social Hymenoptera (bees, wasps, ants and termites) the eggs laid by worker bees are actively killed by others (for a recent overview, see (Ratnieks and Wenseleers 2008).

Social dilemmas, in which individual interests of an individual are opposed to that of the group, have attracted most of the attention in cooperation research. However, in this thesis I will not restrict myself to this specific form of social interaction. Humans are involved in many types of social interaction that are structured differently. For instance, individuals frequently have to coordinate a specific course of action; some outcomes may lead to higher payoffs than others, but it is in the interest of each individual to perform the same behaviour. Examples of such 'coordination games' range from the Stag-Hunt game to driving on either the left or the right side of the road. In other types of interactions, it is beneficial to deviate from the behaviour of others. For instance, in deciding between two patches to forage in, competition can be avoided by choosing the least popular one. Another example of such an 'evasion game', in some situations costly cooperation is the only way to achieve a positive outcome for all involved parties, and forms of specialisation can play a role (*e.g.*, in obtaining a resource, where partners can benefit from trading what they have produced). Each of these social contexts involves forms of cooperation, but each of them has its own underlying structure, calling for specific strategies to reach an efficient outcome.

Social learning and cultural evolution

Cooperation may have been essential in human evolution by allowing us to deal with complex adaptive problems, but our capacity for social learning perhaps contributed even more to our success. The colonisation of the globe by modern humans started from Africa, some 60,000 years ago. Within 20,000 – 30,000 years, humans had spread out across Asia and Australia, with populations living in the arctic (Klein 1989). In the process of colonisation, humans were confronted with a vast array of very

different environments each presenting a unique combination of adaptive problems; newly encountered habitats called for new knowledge, new tools, and new social arrangements in order to survive (Boyd et al. 2011a).

Our ability to adapt to such a wide range of environments is often attributed to our intelligence (Barrett et al. 2007; Pinker 2010), and it is hardly disputable that our ability to make inferences about the world helped to invent new technologies and refine them. This explanation, however, is far from complete. Dramatic examples illustrating that being smart is not enough, are given by expeditions that are stranded in some alien habitat. Despite elaborate efforts and plenty of time to experiment and learn, the expedition members often suffer badly or even die because they lack essential information about the environment, and are unable to make a living on their own. The expedition members that survived such a stranding often survive thanks to the indigenous population taking care of them (Boyd et al. 2011a); similar illustrations are found in (Henrich and McElreath 2003)). This highlights that the human ability to survive in the diverse range of environments that we inhabit critically hinges on adaptive information specific to a certain habitat. Through social learning, populations living in an environment can pass on and accumulate information over generations. This process makes it possible to create tools adapted to deal with specific combinations of adaptive challenges, and to develop technologies of a degree of complexity and sophistication that can never be developed by any individual within his or her life time (Boyd et al. 2011a).

The term social learning refers to a process in which individuals are able to modify their behaviour through observing or interacting with other individuals (Galef Jr 1976; Heyes 1994). This modification of behaviour can take place through a range of different psychological mechanisms (see *e.g.*, (Rendell et al. 2011) for an overview). In this thesis I will focus on social learning by imitation of behaviour performed by others. Just as with cooperation, social learning is by no means unique to humans, and is widely studied in a range of animal species. For example, chimpanzees (*Pan trochodytus*), nine-spined sticklebacks (*Pungitius pungitius*), and even fruit flies (*Drosophila melanogaster*) have been observed to copy behaviour that is performed by conspecifics (Pike and Laland 2010; Battesti et al. 2012; van de Waal et al. 2013). Social learning can be of

great value: by learning from others, individuals can acquire information about important aspects of their physical and social environment such as the location of food and water sources, the avoidance of predators and availability of potential mates (Boyd and Richerson 1985; Rogers 1988; Laland 2004). By gathering information through social learning, an individual can avoid the costs associated with individual (non-social) learning, in which individuals learn by directly interacting with their environment. Moreover, payoffs of behaviour are often subject to noise due to exogenous factors. Such chance effects can make it useful to aggregate information on peer behaviour, so that multiple noisy pieces of information can be turned into a reliable signal that points towards optimal behaviour in a given situation.

Theoretical models suggest that specific strategies of social learning can be adaptive in specific kinds of situations (Laland 2004). Such strategies specify for an individual when to learn from others, from whom to learn, and what kinds of social information to attend to (Laland 2004; Rendell et al. 2011). A set of models suggests that an individual should use social learning when its current behaviour is unproductive, when non-social learning is costly, or when the individual is uncertain about which behaviour fits best in the current circumstances (Boyd and Richerson 1985; Boyd and Richerson 1988; Feldman et al. 1996; Wakano et al. 2004). Another set of models deals with strategies that specify from whom an individual should learn. In some situations, individuals can benefit from attending to the frequency of behaviour in a population, and may be inclined to adopt behaviour that is common (Boyd and Richerson 1985; Godin and Dugatkin 1996; Richerson and Boyd 2004). In other situations, individuals may prefer behaviour that is rare (Griskevicius et al. 2006). In case individuals can estimate the success or performance of their peers, they may even adopt more sophisticated strategies, such as payoff-based learning. Individuals can profit from imitating high-rank or prestigious peers, or copying those peers that do better than themselves (Schlag 1998; Gintis 2000; Henrich and Gil-White 2001). In addition, the content of social information likely influences the probability that traits transmit between individuals. For instance, experimental evidence suggests that information about social relationships is transmitted more frequently and

more accurately than information about individual behaviour or information about the environment (Mesoudi et al. 2006).

The passing of socially learned information between individuals has interesting similarities with the process of the inheritance of genes from parents to offspring. In order to contrast this process with genetic transmission, it has been termed 'cultural transmission'. The analogy between cultural and genetic transmission has been popularised by Dawkins (Dawkins 1976). Views, ideas or beliefs (or 'memes', as Dawkins calls them) can vary within a population, and be transmitted between individuals. From time to time, individuals experiment with their environment, try out new behaviours or invent new techniques. The resulting 'innovations' add to the variation of ideas in the populations, and may be viewed as cultural analogues of mutations. Moreover, cultural variants can be subject to forms of selection when individuals in a population use specific forms of social learning (*e.g.*, conformism or payoff-based learning, or when some variants are more likely to be copied than others due to their content). Because cultural variants are heritable, vary within a population and can be subject to selection, it is possible for cultural traits to evolve over time. In this light, it is not surprising that theories of cultural evolution take models of genetic evolution as the point of departure to analyse how the distribution of culturally inherited traits changes over time (Cavalli-Sforza and Feldman 1981; Boyd and Richerson 1985).

While theoretical research on cultural evolution has lifted off since the 1980's, almost nothing is known about how cultural variants are represented in neural structures of the brains of their bearers. For instance, it is unknown whether cultural transmission involves the particulate inheritance of discrete units of information. It seems that cultural inheritance is in many cases not particulate; think of political beliefs that tend to vary continuously from left wing to right wing, or the width and length of arrow heads that vary on a continuum (Mesoudi 2011b). However, it might be that – in similarity to the genetic inheritance system – phenotypes can vary continuously while the underlying units of transmission are particulate (think of the length or height of an organism as an example). Because of the lack of an empirical basis for the units of inheritance, I will in this thesis abstract from the physical structures

underlying cultural traits and the specifics of their realisation in different individuals. I will take cultural variants as behavioural traits that can be transmitted between individuals through social learning, and are able to affect individual behaviour.

The direction and outcome of cultural evolution depend on the social learning strategies that individuals in a population use. For instance, when individuals in a population tend to learn from more successful peers (*e.g.*, peers that are more efficient in dealing with their environment), cultural evolution can have dynamics that are similar to those of genetic evolution by natural selection. In genetic evolution, natural selection leads to the preferential transmission of genetic variants that enhance the performance (or 'fitness') of their bearers. Similarly, when individuals tend to adopt traits by payoff-based learning, variants that enhance the performance of individuals are transmitted more frequently. As a result, adaptive behaviour can spread through a population. Alternatively, when individuals tend to follow the majority in determining their behaviour, common traits will tend to become even more common, whereas rare traits tend to decrease in frequency even more. The outcome of cultural evolution under such conformist learning thus depends on the initial conditions present in the population. The differences and similarities between genetic and cultural evolution are discussed further below.

In this thesis I aim to contribute to our understanding of cultural evolution, both by developing models and by conducting experiments that test the assumptions and predictions of these models. First, I systematically map out the consequences of various forms of social learning – such as following a leader, or following the majority – for the direction and outcome of cultural evolution in a number of interaction contexts (Chapters 2 and 3). Second, I investigate what forms of social learning people actually use in such situations, and whether individuals are consistent in their use of their social learning strategies (Chapters 4 and 5). Third, I assess what historical information about social interactions people use when cooperating with interaction partners (Chapters 6 and 7). The following section of this introduction will further introduce the connection between social learning and human cooperation, as well as their interplay in shaping human behaviour.

Social learning and human cooperation

So far I have discussed social learning and human cooperation separately. However, these two processes are intimately linked. Sharing valuable information through communication presents a form of cooperation between individuals. Many forms of social learning depend on people's willingness to take time to teach, demonstrate specific techniques and transmit social norms to others. Conversely, social learning can have a great impact on human cooperation. An influential view holds that when behaviours transmit culturally, rather than genetically, cooperation is more likely to evolve (Boyd and Richerson 1985; Henrich 2004). The advocates of this view use models of cultural transmission that help to explain the high levels of cooperation between humans, as well as the diversity in norms of cooperation between cultures (Boyd and Richerson 1985; Henrich and Boyd 1998; Henrich 2004; Richerson and Boyd 2004). This strand of models is built upon the insight that, despite the similarities between genetic and cultural transmission, the laws of inheritance between them differ in a number of ways. The most important difference is that cultural evolution is argued to be based not only on payoff-based transmission (which would lead to dynamics similar to those of genetic evolution), but also largely on other forms of social learning, and conformism in particular. As a consequence, genetic and cultural evolution can lead to very distinct outcomes.

To see how these different outcomes are reached in contexts of human cooperation, consider a population of individuals involved in a social dilemma. For clarity of exposition, let us focus on a simple scenario in which individuals have two behavioural options: they can either choose whether or not to contribute to a joint project that benefits all group members (for instance, helping to defend a tribal village from an outside threat). In this interaction context, all individuals in a group obtain higher payoffs when they all contribute (where all members cooperate and benefit from a safe village), compared to a situation in which all individuals defect (where none of the members protects the village, resulting in frequent raids by bandits). Individually, however, the highest payoffs are acquired by refraining from contributing to the defence, reaping the benefits of safety without paying the costs of risking injuries. In other words, collective interests and individual interests oppose each

other. Now suppose that an individual's tendency to cooperate is governed by a gene with two alleles: 'cooperate' and 'defect'. Everything else being equal, natural selection would predict that the 'defect' allele increases in frequency due to its higher payoffs, and that cooperators will eventually go extinct. Similarly, if the tendency to cooperate is governed by a socially learned trait, and individuals use payoff-based learning to update this tendency, cooperation is likely to decrease. When individuals use other forms of social learning, however, cooperation might be stable. For instance, when a tribe mainly consists of cooperative individuals, conformist social learning is likely to lead to a situation where all individuals cooperate.

The argument above sketches a scenario in which cultural evolution can lead to cooperation within one group. The advocates of this argument often invoke 'cultural group selection' as a mechanism of how cooperation can spread throughout a larger metapopulation structured into different groups (Henrich and Boyd 1998; Boyd et al. 2003; Henrich 2004). As already suggested by Darwin, selection is not by definition limited to the level of individuals, and anthropological evidence suggests that fierce competition between tribes of humans was quite common in our evolutionary history (Darwin 1871; Soltis et al. 1995). It is plausible that high degrees of cooperation enhanced the ability of prehistoric tribes to maintain superior systems of subsistence, sustain healthier warriors, and have a bigger army with superior war technologies. As a consequence, tribes of selfish individuals are likely to lose wars and to be supplanted by more cooperative tribes, or be taken up in these tribes and, over time, adopt their cultural traits. In Chapters 2 and 3 of this thesis, I discuss to what extent cultural evolution can lead to stable cooperation within one group and how specific forms of between-group selection can lead to the spread of cooperation on a larger scale.

In a social dilemma situation, within-group selection (through payoff-based learning) favours defection, whereas between-group selection favours cooperation. So, at these two levels, selection works in opposite directions. The net direction of change depends on the relative strength of the selective forces, as well as on the amount of variation on each level (Price 1970). For group selection to facilitate the evolution of cooperation, the behavioural variation between groups should be large compared to

the variation within groups. Since conformist social learning tends to reduce variation within groups, it can promote the cultural evolution of cooperation by group selection (Boyd and Richerson 1985; Henrich and Boyd 1998; Henrich and Boyd 2001; Bowles 2001; Richerson and Boyd 2004). This argument suggests that in principle, specific forms of social learning (conformist learning in particular) may widen the range of conditions under which cooperation is likely to evolve. However, a range of studies taking a different modelling approach show that the cultural evolution of cooperation strongly depends on specific assumptions on the demographics of the population (*e.g.*, the way the population is structured, and rates of migration between groups) and the way that cooperation can spread between groups (*e.g.*, by differential migration or by replacement of groups; (Lehmann and Feldman 2008; Lehmann et al. 2008a; Lehmann et al. 2008b). Moreover, these authors conclude that conformism can hinder rather than promote the cultural evolution of cooperation. This strongly contrasting conclusion led to a hot debate which is still ongoing in the literature (Boyd et al. 2011b). This issue is taken up in Chapter 3 of this thesis, in which I offer a systematic comparison of the effects of various social learning strategies and forms of group selection within one modelling framework.

Perhaps the largest obstacle for progress in understanding of cultural evolution is fact that model assumptions on how traits transmit between individuals generally lack a solid empirical justification. Assumptions on social learning strategies – and conformism in particular – are often based on theoretical considerations (*e.g.*, (Boyd and Richerson 1985; Henrich and Boyd 1998; Richerson and Boyd 2004). Elsewhere, assumptions on social learning are justified by appealing to their ability to explain a range of social phenomena, such as the persistence of maladaptive social norms (Bowles 2001; Henrich 2004). Most of the empirical evidence supporting these models derives from social psychological experiments in which students are confronted with decision making contexts different from cooperation or coordination problems, or not specifically targeted to measure social learning strategies (*e.g.*, (Asch 1956; Insko et al. 1985; Wit 1999).

Over the last decade, however, a body of experimental data has started to accumulate shedding light on the way people learn from each other

(McElreath et al. 2005; Efferson et al. 2007; McElreath et al. 2008; Efferson et al. 2008; Mesoudi 2008; Mesoudi et al. 2008; Traulsen et al. 2010; Mesoudi 2011a). These experiments on human social learning strategies are the first steps towards establishing an empirically grounded theory of cultural evolution. In a typical experiment, people are invited into a computer laboratory and repeatedly asked to make decisions under controlled conditions. Participants can choose between a number of different options, each of which is associated with a certain payoff. Payoffs are noisy, so it can be useful to collect social information to find out which behaviour leads to the best outcome. In the controlled circumstances of a decision making laboratory, aspects of social learning strategies can be estimated by documenting which social information people attend to before making their decisions, and how such information affects the choices they make. The picture of human social learning strategies emerging from these experiments shows that payoffs as well as frequencies of peers' behaviour play a role in human social learning. The growing body of evidence is, however largely confined to situations in which individuals have to deal with a single non-social problem, such as finding out which technology works best in a newly encountered environment. However, to understand the role of cultural transmission in the evolution of human cooperation, it is crucial to know how people learn from each other in social situations where payoffs of behaviour depend on what others are doing. To the best of my knowledge, almost no empirical evidence exists for human social learning strategies in such situations. By studying experimentally how people learn from each other in a range of social contexts, I aim to contribute to the empirical groundwork for theories of the cultural evolution of cooperation.

In this thesis I study the interplay between human cooperation and social learning. By theoretical investigation of the interactive effects of different forms of social learning and group selection I aim to shed light on the extent to which cultural evolution can lead to the evolution of cooperation (Chapters 2 and 3). Moreover, I provide experimental evidence on human social learning strategies in a variety of social contexts (Chapters 4 and 5). Additional analyses of strategic diversity in cooperative behaviour – in particular, with respect to reciprocal helping behaviour – can be found towards the end of this thesis (Chapters 6 and 7).

This thesis

In Chapters 2 and 3 of this thesis I explore how various forms of social learning affect the distribution of behavioural social strategies over time. These two chapters provide the general theoretical framework for this thesis. The models presented in this part are conceptual in nature; their purpose is not to represent a complete and realistic image of how social learning works in humans. Rather, they single out factors that may be of importance in causing cultural change – and study their impact on the direction and outcome of cultural evolution.

In **Chapter 2**, I offer a systematic survey of how two prominent forms of social learning – conformism and payoff-based learning – affect the outcome of cultural evolution in a number of interaction contexts. I offer a set of game theoretic models in which individuals interact in a range of different social settings, in each of which the payoffs of an individual's behaviour depend on the behaviour of others in a different way. In determining their behaviour, individuals can occasionally update their behavioural strategy by imitating other members of their group. This updating occurs based on the payoffs or the frequency of a certain behaviour. In each of these games, I study how the direction and outcome of cultural evolution is affected by the relative frequency of either social learning rule. I conclude this chapter with an analysis of the effect of conformism on the evolution of cooperation by cultural group selection. I do this by focusing on a stochastic model of cultural evolution in a social dilemma in a group-structured population, allowing individuals to occasionally update their behaviour by comparing their payoffs to members from other groups, thereby introducing the group as a level of selection to the system.

Chapter 3 takes up the latter issue and digs deeper into the dynamics of cultural group selection in populations of individuals that face a social dilemma. Using individual-based simulations, I further investigate how various forms of social learning affect the outcome of cultural evolution. More specifically, I intend to help resolve a hot issue in the current literature on whether, and to what extent, social learning promotes the evolution of cooperation.

Chapters 4 and 5 of this thesis aim to support theories of cultural evolution and cooperation with empirical evidence. I study human social learning strategies and cooperative interactions – as well as the interplay between these two – using techniques borrowed from experimental economics. I present the results of a set of decision making experiments on human social learning strategies. These experiments aim to provide empirical support for theories of cultural evolution by studying the mechanisms of cultural transmission. Specifically, **Chapter 4** focuses on the use of social information in a range of interaction contexts. I study how individuals make use of social information in making decisions under uncertainty, and focus on the consistency in their individual social learning strategies across interaction contexts. **Chapter 5** complements these findings by investigating the extent to which these strategies are stable over time. In addition, I investigate potential links between social learning strategies and the outcome of group interactions in a social dilemma.

Chapters 6 and 7 focus on the issue of how reciprocity affects the outcome of cooperative interactions. When deciding to help someone, people often take into account this person's history of cooperative behaviour towards themselves through direct experience (direct reciprocity), and towards others through reputations (indirect reciprocity). In **Chapter 6** I investigate the interplay of direct and indirect reciprocity in an experimental setting. These mechanisms *together* governed human cooperation for most of our evolutionary history, but thus far, the effects of direct and indirect reciprocity were always studied in isolation. I describe the results of a decision making experiment shedding light on the concerted action of personal experience and reputations on the dynamics of cooperation in groups. **Chapter 7** zooms in on the dynamics of indirect reciprocity. I present the results of an experiment designed to study how individuals condition their cooperation decisions on past behaviour of their interaction partners. In particular, our setup allows individuals to observe the information that their interaction partners had, when deciding to help others. By analysing how individuals react to such second-order information, I obtain insight into the extent to which individuals consider the 'righteousness' of their partners, and the

consequences of individual variation in strategies on the outcome of social interactions.

Chapter 8 summarises the main results presented in this thesis, and places these findings into a broader context. I conclude by delineating issues that remained unaddressed here, and describe potentially interesting avenues for future research.

Chapter two

Effects of conformism on the cultural evolution of social behaviour

Lucas Molleman
Ido Pen
Franz J. Weissing

PLoS One 8 (7): e68153. doi:10.1371/journal.pone.0068153

Abstract

Models of cultural evolution study how the distribution of cultural traits changes over time. The dynamics of cultural evolution strongly depends on the way these traits are transmitted between individuals by social learning. Two prominent forms of social learning are payoff-based learning (imitating others that have higher payoffs) and conformist learning (imitating locally common behaviours). How payoff-based and conformist learning affect the cultural evolution of cooperation is currently a matter of lively debate, but few studies systematically analyse the interplay of these forms of social learning. Here we perform such a study by investigating how the interaction of payoff-based and conformist learning affects the outcome of cultural evolution in three social contexts. First, we develop a simple argument that provides insights into how the outcome of cultural evolution will change when more and more conformist learning is added to payoff-based learning. In a social dilemma (*e.g.*, a Prisoner's Dilemma), conformism can turn cooperation into a stable equilibrium; in an evasion game (*e.g.*, a Hawk-Dove game or a Snowdrift game) conformism tends to destabilize the polymorphic equilibrium; and in a coordination game (*e.g.*, a Stag Hunt game), conformism changes the basin of attraction of the two equilibria. Second, we analyse a stochastic event-based model, revealing that conformism increases the speed of cultural evolution towards pure equilibria. Individual-based simulations as well as the analysis of the diffusion approximation of the stochastic model by and large confirm our findings. Third, we investigate the effect of an increasing degree of conformism on cultural group selection in a group-structured population. We conclude that, in contrast to statements in the literature, conformism hinders rather than promotes the evolution of cooperation.

Introduction

Social learning enables humans to survive in a broad array of different habitats across the planet. By learning from their peers, individuals can rapidly acquire adaptive information about which behaviour is optimal under a variety of environmental conditions. Models of cultural evolution use insights from theories of genetic evolution to study how cultural variants, such as ideas and beliefs, spread through populations of individuals by social learning. Social learning based on imitating the behaviour of successful individuals can lead to an evolutionary dynamic similar to the spread of alleles under natural selection, whereas learning by adopting behaviours from others more randomly leads to a process resembling genetic drift.

Models of cultural evolution have to be adapted to the specific mechanisms by which cultural traits transmit between individuals. Traits can be transmitted not only vertically from parents to offspring, but in a range of different ways. For instance, traditional hunters may learn from their parents a social norm to share hunting revenues, and may learn the optimal design of an arrow from their fellow hunters. How humans learn from each other is a topic of extensive theoretical and empirical research (for a recent overview see Rendell et al. (2011)), and various specific forms of social learning (termed 'social learning strategies' (Laland 2004) or 'learning biases' (Boyd and Richerson 1985)) have been studied as to how they affect the spread of cultural traits through populations. Two forms of social learning received particular attention: conformism and payoff-based learning.

When individuals can evaluate the payoffs of the behavioural strategies of others, the preferential imitation of high-payoff individuals can lead to the rapid spread of adaptive behaviours in a population (Henrich and Gil-White 2001). However, such payoff-based learning is not always feasible. Getting insights in the payoffs received by others is not always straightforward, especially for newcomers in a population. In cases like this, imitating the majority (conformism) can be a good alternative form of social learning, in particular if the success of cultural traits strongly depends on the local circumstances (Boyd and Richerson 1985; Henrich and Boyd 1998).

The role of conformism in cultural evolution has recently become the matter of considerable debate. In the context of a social dilemma, payoff-based learning will tend to inhibit the spread of cooperation because defectors obtain higher payoffs by reaping the benefits of cooperation without paying the costs. Theory suggest that when payoff-based learning is complemented by other forms of social learning, the dynamics of cultural evolution can be strongly affected. For instance, adding random learning to payoff-based learning can facilitate the rapid solution of coordination problems (Vilone et al. 2012), and conformism can stabilise cooperative equilibria under specific conditions (Boyd and Richerson 1985; Henrich and Boyd 1998; Henrich and Boyd 2001; Nowak et al. 2004; Skyrms 2005; Traulsen et al. 2006; Andrés Guzmán et al. 2007; Peña et al. 2009). Moreover, conformism can homogenise groups internally, thereby augmenting the relative amount of variation *between* groups (Henrich and Boyd 1998). This decreases the scope for selection within groups (*i.e.*, payoff-based transmission disfavouring cooperation), and increases the potential role of ‘cultural group selection’. In group-structured populations, cooperation can spread when groups of cooperators have some advantage over groups of defectors. This advantage can manifest itself in a number of different ways; cooperative groups may send out more migrants, grow to larger sizes, or replace other groups (*e.g.*, Boyd et al. 2003; Henrich 2004; Traulsen and Nowak 2006; Andrés Guzmán et al. 2007). Selection at the group level may also occur when individuals occasionally learn from members of other groups that perform better (Henrich and Gil-White 2001; Boyd and Richerson 2002; Lehmann et al. 2008; Lehmann and Feldman 2008). Through such a process, cooperation can be promoted since individuals in cooperative groups have higher payoffs than individuals in groups of non-cooperators.

Experiments from psychology and behavioural economics suggest that humans indeed use both conformist and payoff-based learning in determining their behaviour (Asch 1956; McElreath et al. 2005; Efferson et al. 2008; Traulsen et al. 2010). When individuals are allowed to use both conformist and payoff-based learning, experimental evidence suggests that cultural traits can spread through a *mixture* of these two forms of social learning (McElreath et al. 2005; McElreath et al. 2008; Mesoudi 2011). This raises the question of how the interplay of

conformist and payoff-based learning affects the spread of cultural traits through a population.

First, we develop a simple argument to delineate how the direction of cultural change is affected by the relative rate of conformist and payoff-based learning. This will give us an intuitive insight in the effects of conformism in various contexts of social interaction. Second, we construct stochastic models that allow us to follow the spread of culturally transmitted behaviours in the course of time in a finite population. These models allow us to quantify how the relative degree of conformism (as opposed to payoff-based learning) affects the success of social strategies in reaching fixation. Third, we examine how cooperation can spread in a group-structured population by means of cultural group selection. With this model, we investigate whether conformism tends to promote – as often claimed in the literature – or hamper the spread of cooperation in populations that are structured into groups of finite size.

Analysis and results

2.1. Model structure

We consider a population in which individuals are involved in social interactions. Individuals have a culturally acquired strategy that determines their behaviour in these interactions. We consider two variants of this behaviour (*A* or *B*). An individual has the inclination of playing either *A* or *B*, but this inclination can change over the course of time due to social learning. Learning is either based on payoffs (individuals tend to imitate successful individuals) or on conformism (individuals tend to imitate the majority of the population). The relative frequency γ of these two forms of social learning is the key parameter of interest. The value of γ ranges from 0 to 1. If $\gamma=0$, all learning is payoff-based; if $\gamma=1$, all learning is based on conformism. If $0 < \gamma < 1$, individuals use a mixture of these two forms of social learning. We assume that all individuals use the same mixture of conformist and payoff-based learning.

Individuals acquire payoffs by social interaction with others in their group. Payoffs depend linearly on the frequency p of *A*-individuals (and $1-p$).

p of B -individuals), and are calculated using the payoff matrix

$G = \begin{pmatrix} a & b \\ c & d \end{pmatrix}$. The payoff of A -individuals is

$$\pi_A = p \cdot a + (1-p) \cdot b \quad (2.1a)$$

and the payoff of B -individuals is

$$\pi_B = p \cdot c + (1-p) \cdot d. \quad (2.1b)$$

There are three strategically different classes of games with two pure strategies, and we consider the evolutionary dynamics in each of these 'interaction contexts'. In the first class of games, one of the pure strategies (say B) is dominant over the other: $a < c$ and $b < d$. In the special case where $a > d$, this is a *social dilemma*. Collective interests are opposed to individual interests: when all individuals exhibit behaviour A ('cooperate'), payoffs are higher than when all individuals exhibit behaviour B ('defect'). Individually, however, B yields higher payoffs than A , irrespective of what others are doing. Second, we consider the class of *coordination games*, which are characterized by $a > c$ and $b < d$. In this case, the payoff of a pure strategy increases with the number of individuals using this strategy. In a coordination game, both pure strategies are Nash equilibrium strategies. In addition, there is a (dynamically unstable) mixed-strategy equilibrium at

$$p^* = \frac{d-b}{a-b-c+d}. \quad (2.2)$$

Third, we consider the class of *evasion games*, where $a < c$ and $b > d$. Now the relative payoff of each pure strategy decreases with the frequency of this strategy in the population. In an evasion game, none of the pure strategies is a Nash equilibrium; instead an evasion game has a unique mixed-strategy equilibrium, which is given by eq. (2.2). The Hawk-Dove game and the Snowdrift game are prominent examples of evasion games.

2.2. Representation of conformism by a coordination game

Before studying the dynamics of cultural evolution in finite populations, we aim to get some intuition on how conformism might change the direction of cultural evolution. To this end, we represent conformism-based learning by payoff-based learning in the coordination game with payoff matrix $K = \begin{pmatrix} s & 0 \\ 0 & s \end{pmatrix}$. In this coordination game, the payoff to each

pure strategy is proportional to its frequency in the population: $\pi_A = p \cdot s$ and $\pi_B = (1-p) \cdot s$, where s is positive. Accordingly, the pure strategy with highest frequency is favoured by payoff-based learning, just as it is in case of conformism-based learning. Based on these considerations, a mixture of payoff-based learning (characterised by matrix G) and conformist learning (characterised by matrix K) can be described by the combined matrix

$$M(\gamma) = (1-\gamma) \cdot G + \gamma \cdot K = \begin{pmatrix} (1-\gamma)a + \gamma s & (1-\gamma)b \\ (1-\gamma)c & (1-\gamma)d + \gamma s \end{pmatrix}. \quad (2.3)$$

It is now straightforward to characterise the expected direction of cultural change as a function of our key parameter γ by determining the Nash equilibrium strategies of the matrix game $M(\gamma)$. This can be done with standard methods (Osborne and Rubinstein 1994): Pure strategy A is a

Nash equilibrium if $(1-\gamma)a + \gamma s > (1-\gamma)c$ or equivalently $\frac{\gamma}{1-\gamma} s > c - a$.

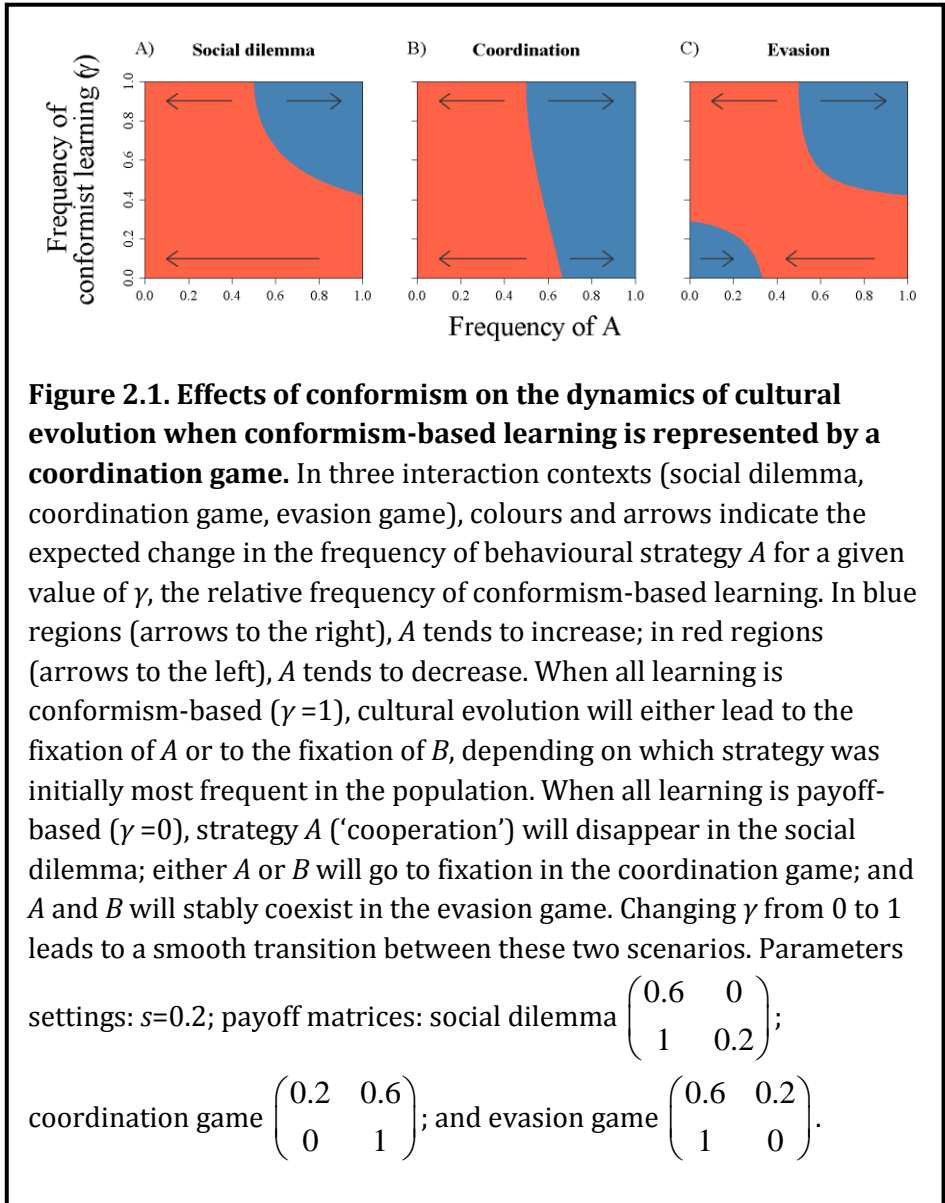
Similarly, pure strategy B is a Nash equilibrium when $\frac{\gamma}{1-\gamma} s > b - d$. Both

inequalities are more easily satisfied for larger values of s or γ and will always hold if γ approaches 1. When both inequalities are reversed, $M(\gamma)$ has a (dynamically stable) mixed Nash equilibrium at

$$p^* = \frac{(1-\gamma)(d-b) + \gamma s}{(1-\gamma)(a-b-c+d) + 2\gamma s}. \quad (2.4)$$

For each of the three interaction contexts, Figure 2.1 illustrates how the dynamics of cultural evolution (increase or decrease in the frequency of pure strategy A) changes with the frequency of conformist learning γ .

First, consider the extreme case $\gamma = 0$, at the bottom of the three panels of Figure 2.1. Here, all learning occurs on the basis of the payoffs in matrix G . In the *social dilemma*, A (cooperate) is disfavoured by payoff-based learning, and cultural evolution will lead to a decrease of the frequency of A (Figure 2.1A, bottom arrow to the left) and convergence to the sole Nash equilibrium $p^* = 0$. The *coordination game* has two pure-strategy Nash equilibria ($p^* = 0$ and $p^* = 1$) that are separated by the dynamically unstable mixed-strategy Nash equilibrium (eq. 2.2). The two arrows at the bottom of Figure 2.1B indicate that cultural evolution will either lead to the fixation of A or to the fixation of B , and that the outcome depends on initial conditions. In the *evasion game*, the arrows at the bottom of Figure 2.1C indicate that the system will converge to the mixed-strategy Nash equilibrium (eq. 2.2), where A and B stably coexist. Next consider the other extreme $\gamma = 1$, where all learning is conformism based, *i.e.* governed by matrix K (top of the three panels in Figure 2.1). Now the expected direction of change is identical for each of the three interaction contexts: since K is a coordination game, the two pure strategies are Nash equilibria and cultural evolution will either lead to the fixation of A or to the fixation of B , depending on initial conditions (top arrows in all three panels). Due to conformist learning, the strategy that is initially more abundant is most likely to spread to fixation.



For intermediate frequencies of conformism $0 < \gamma < 1$, we observe a gradual shift between the two extreme cases $\gamma = 0$ and $\gamma = 1$. In the *social dilemma* (Figure 2.1A), cooperation becomes a stable equilibrium as soon as conformism is sufficiently strong (i.e. if $\frac{\gamma}{1-\gamma} s > c - a$): when a group

mainly consists of cooperators, conformist learning leads to the maintenance and fixation of this most abundant strategy, despite of its payoff disadvantage. In the *coordination game* (Figure 2.1B), the direction of change remains qualitatively unchanged, but the unstable Nash equilibrium separating the basins of attraction shifts from eq. (2.2) to $p^* = 0.5$, the mixed-strategy Nash equilibrium of the coordination game K . In the *evasion game* (Figure 2.1C), A and B coexist at equilibrium when learning is mainly based on payoffs (small γ). High frequencies of conformism γ decrease the scope for this coexistence. When conformism occurs at a sufficiently high frequency (moving towards the top of the panel), the direction of expected change is reversed, and the two monomorphic equilibria become stable.

2.3. A stochastic model for cultural evolution in finite populations

The approach taken above provides an intuitive understanding of the role of conformism in various types of interaction contexts. However, it is not clear whether, and to what extent, the features of conformism-based learning are captured by a coordination game. We therefore developed a dynamic model for cultural evolution where conformism is represented in a more mechanistic way. Cultural evolution takes place in finite (and often small) populations, where chance events may play an important role. Accordingly, we consider a stochastic model for cultural evolution in a finite population of fixed size n . Our model is event-based, where an ‘event’ corresponds to a potential change in strategy by one population member. In each time step (i.e. when an event occurs) two individuals are chosen at random from the population, and one of them is allowed to update its behavioural strategy (A or B) by learning from the other. Updating occurs either through conformist or payoff-based learning. For each state of the population (i.e. each possible frequency of A -strategists), we calculate the probability that an A -individual switches to B , and that a B -individual switches to A . At the population level, each such switch

corresponds to a decrease or to an increase of the number of A -individuals by one. Since we neglect the spontaneous emergence of A - or B -strategists (the cultural equivalent of genetic mutations), the stochastic process will eventually lead to an ‘absorbing state’, where all individuals have adopted either of the two strategies. To assess the effects of conformist learning on the outcome of cultural evolution, we evaluate how our key parameter γ affects the fixation probability and waiting time to fixation for each of the two behavioural strategies (Karlin and Taylor 1975).

Let i be the number of A -strategists in the population, and let T_i^+ and T_i^- denote the probability of gaining resp. losing one A -individual. We model the switching dynamics by using pairwise comparison (*cf.* Traulsen et al. 2006). When an event occurs, two individuals are chosen at random from the population. A change in strategy can only take place when these two individuals have opposite strategies, which occurs with probability $\frac{i(n-i)}{n(n-1)}$. In state i , switching from A to B (probability T_i^+) and from B to A (probability T_i^-) occurs with probabilities:

$$T_i^\pm = \frac{i(n-i)}{n(n-1)} \cdot \left[(1-\gamma) \cdot P_i^\pm + \gamma \cdot C_i^\pm \right] \quad (2.5)$$

where C and P denote the probabilities of switching due to conformist and payoff-based learning, respectively. For both forms of social learning, we specify the switching probabilities as a logistic function of the differences in payoffs (in case of payoff-based learning) or frequencies (in case of conformist learning) between strategies A and B . In state i , payoff-based switching from B to A (P_i^+), and from A to B (P_i^-) occurs with probabilities

$$P_i^+ = \Pr(B \rightarrow A | i) = \left[1 + \exp(-\beta_P \cdot (\pi_B(i) - \pi_A(i))) \right]^{-1} \quad (2.6a)$$

$$P_i^- = \Pr(A \rightarrow B | i) = \left[1 + \exp(-\beta_P \cdot (\pi_A(i) - \pi_B(i))) \right]^{-1} \quad (2.6b)$$

where $\pi_A(i)$ and $\pi_B(i)$ (see equation (2.1)) refer to the payoffs of strategies A and B in state i , respectively. Parameter β_p quantifies the strength and direction of the relation between the payoff difference and the probability of switching. When $\beta_p=0$, payoff-based learning is not biased in any particular direction, and is expected to lead to dynamics similar to genetic drift. When β_p is large, payoff-based learning is strongly biased, favouring the spread of strategies with the highest payoff (see Figure S2.1 in the Appendix for an illustration).

Conformist learning is represented in a similar way. In state i , conformist switching occurs with probabilities

$$C_i^+ = \Pr(B \rightarrow A | i) = \left[1 + \exp(-\beta_c \cdot (f_B(i) - f_A(i)))\right]^{-1} \quad (2.7a)$$

$$C_i^- = \Pr(A \rightarrow B | i) = \left[1 + \exp(-\beta_c \cdot (f_A(i) - f_B(i)))\right]^{-1} \quad (2.7b)$$

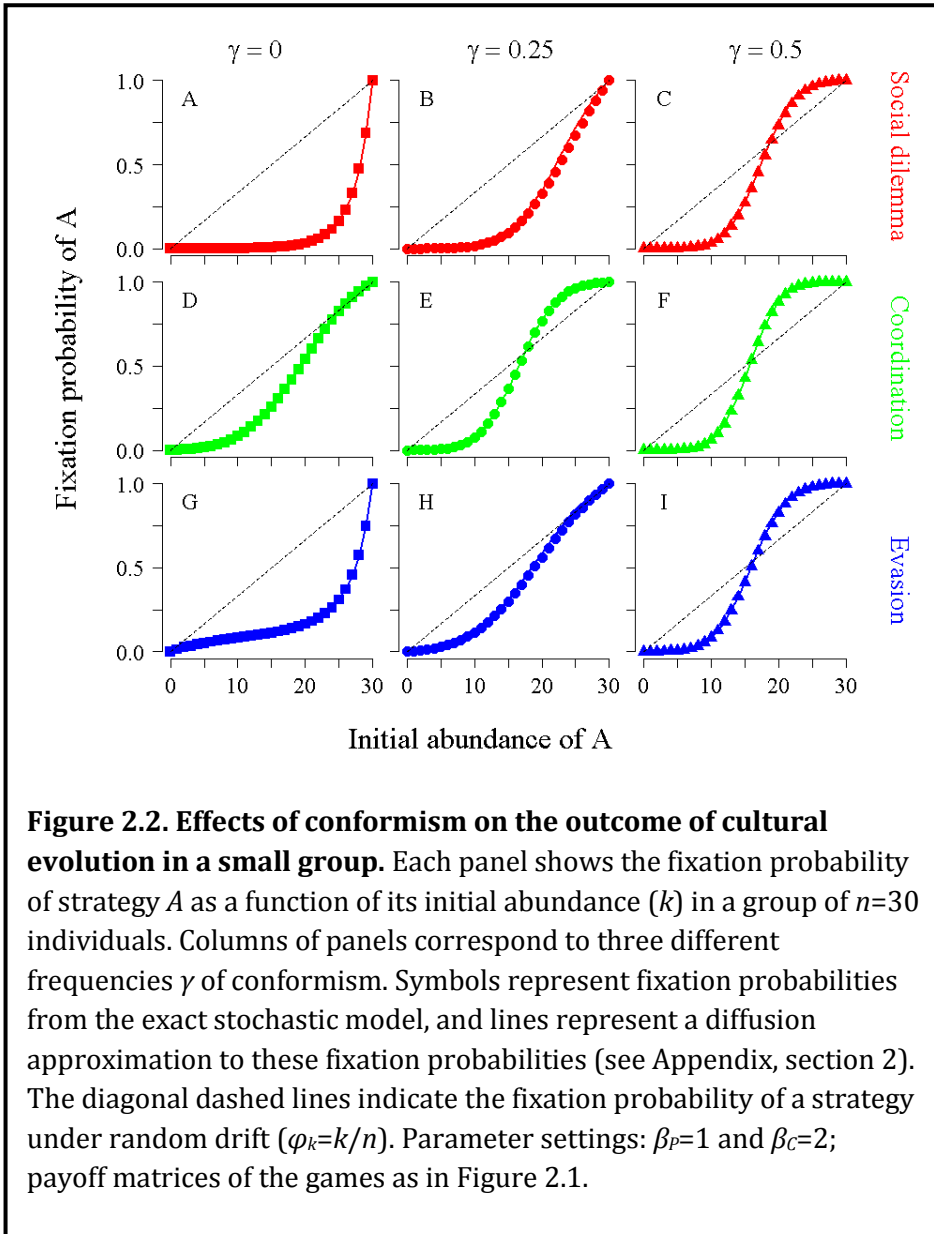
where $f_A(i) = \frac{i}{n}$ and $f_B(i) = \frac{n-i}{n}$ refer to the relative frequencies of strategies A and B , respectively. Parameter β_c quantifies the strength and direction of frequency-based social learning. When $\beta_c=0$, such learning is not biased in any particular direction; when β_c is large, individuals are strongly inclined to adopt the more frequent strategy in the population (see Figure S2.1).

Now we have specified T^+ and T^- for all states i of the population, we can use standard methods (Karlin and Taylor 1975) to calculate fixation probabilities φ and waiting times to fixation, for various initial abundances of strategy A and B , as a function of the relative frequency of conformism γ . In the Appendix, section 2.2, we show how explicit equations for the fixation probabilities can be derived on the basis of a diffusion approximation of the stochastic model.

We assess the influence of the social learning rules on the outcome of cultural evolution by comparing the fixation probabilities φ_k of A for a given initial abundance k of this strategy to that of a ‘neutral’ process, where all switching occurs randomly. It is well known that in the latter

case φ_k equals the initial frequency k/n of A (Karlin and Taylor 1975). Panels A to C in Figure 2.2 illustrate the effects of conformism on the fixation probability of A (cooperate) in a *social dilemma*. When all switching occurs on the basis of payoffs ($\gamma=0$; 2.2A), cooperation is always disfavoured. Accordingly, fixation of A is very unlikely unless the initial abundance of A is relatively large. In fact, all fixation probabilities (red dots) are below the diagonal, indicating that fixation of A is for all values of k less likely than ‘neutral’ updating. If social learning is partly based on conformism ($\gamma=0.25$, 2.2B; $\gamma=0.5$, 2.2C), the fixation curve becomes more S-shaped. Strategy A is increasingly likely to fixate when its initial abundance is high, which agrees with our earlier findings (Figure 2.1A).

In the *coordination game* without conformism (Figure 2.2D), payoff differences near the A equilibrium are relatively small. This means that switching from A to B can frequently occur, despite the fact that strategy B yields lower payoffs. When the state of the group is close to the unstable equilibrium, such stochastic events can tip the group into the basin of attraction of B . As conformism increases in frequency, the pure A equilibrium tends to be more stable; the fixation probability rises above the diagonal k/n line when the initial abundance of A is high (Figure 2.2E, F). This finding is again in line with our earlier results (Figure 2.1B).



In the *evasion game*, payoff-based learning tends to favour the spread of rare strategies. When conformist learning is absent ($\gamma=0$, Figure 2.2G), fixation of B is often more likely. The polymorphic equilibrium is located at $n/3$, and a group is expected to spend most of the time close to this

equilibrium. Since this state is closer to the absorbing state where all individuals play strategy B (relative to the other absorbing state, where all individuals play strategy A), stochastic events will more likely lead to fixation of B rather than fixation of A . When conformist learning occurs at higher frequencies (Figure 2.2H, I), cultural evolution tends to lead to fixation of the strategy that was more abundant initially; again, this is in line with our earlier findings (Figure 2.1C).

Conformism also affects the time it takes until a strategy fixates in the *evasion game* (Figure 2.3). In an evasion game, each behavioural strategy has higher payoff when rare. Accordingly, payoff-based learning causes a group to spend a lot of time in polymorphic states before it reaches one of the absorbing states. Increasing the frequency γ of conformist learning has two effects: first, the frequency of payoff-based switching decreases, which hampers the spread of rare strategies, thereby destabilising the coexistence equilibrium. Second, conformist switching accelerates fixation, because individuals preferentially adopt common strategies.

By breaking down polymorphism in the evasion game, conformism affects the average payoffs of the group members. A simple calculation shows that the average payoff at the mixed-strategy Nash equilibrium (eq. 2.2) of an evasion game is given by $\frac{ad - bc}{a - b - c + d}$. This can be considerably higher or lower than the payoff a in case of fixation of behavioural strategy A or the payoff d in case of fixation of B .

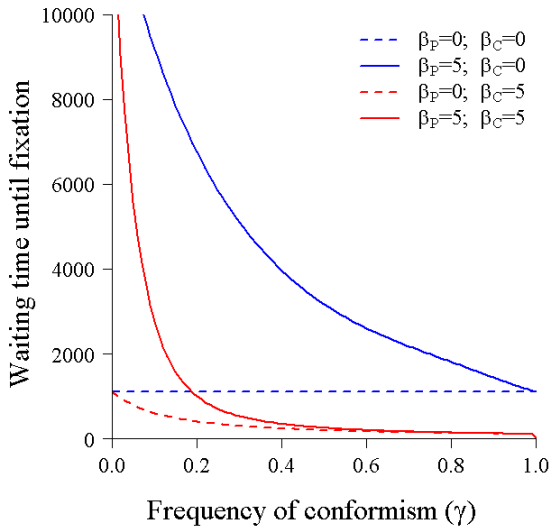


Figure 2.3. Effects of conformism on the persistence of behavioural polymorphism in an evasion game.

Lines represent the expected number of updating events (*i.e.* the ‘waiting time’) until a group fixates in either *A* or *B*, as a function of the frequency of conformist updating γ . The dashed blue line indicates waiting times when all updating occurs randomly ($\beta_p=0$; $\beta_c=0$). The dashed red line reflects waiting times when conformism occurs at rate γ , and payoff-based learning is absent ($\beta_p=0$; $\beta_c=1$; random updating occurs at frequency $1-\gamma$). The solid blue line reflects waiting times when payoff-based learning occurs at frequency $1-\gamma$, and conformism is absent ($\beta_p=1$; $\beta_c=0$; random updating occurs at frequency γ). The solid red line represents waiting times in the full model, where payoff-based learning is complemented by conformism at rate γ ($\beta_p=1$; $\beta_c=1$). Groups were of size $n=30$ and initialised at the coexistence equilibrium $p=p^*$ (10 *A*-individuals). Payoff matrix as in Figure 2.1.

2.4. Evolution of cooperation by cultural group selection

Finally, we consider a multilevel scenario in which a metapopulation is subdivided into m groups of size n . Within groups, individuals face a social dilemma. As described above, payoff-based learning within groups tends to disfavour cooperation. Once in a while, individuals from two *different* groups are paired for updating by comparing their payoffs. This reflects a scenario where individuals occasionally copy behaviours from groups that are performing well. It might be that individuals from other groups are considered to be healthier, or have more wealth. We assume that conformism does not play a role in between-group updating (*i.e.* conformism is a strictly *local* social learning rule, allowing to cope with local conditions). Further, in our model, updating outside the group occurs at a much lower rate than within-group updating: in between two outside-group updating events. This implies that when a new strategy is newly introduced into a group, this strategy will either have gone locally extinct, or reached fixation before the next between-group event occurs. This ‘separation of time scales’ allows us to calculate the probability of fixation of a ‘cooperative’ A strategy in the metapopulation, by tracking the number of groups in the cooperative state. In each time step of this group level process, the abundance of cooperator groups can go up by one or go down by one, or can stay the same.

When two individuals from different groups are chosen from the population, switching probabilities are defined analogously to updating within groups. Switching probabilities depend on the payoff difference between groups where cooperation is fixated and groups where defection is fixated. This payoff difference is given by $a - d$. For this between-group process, we again use logistic functions to specify the relationship between payoff differences and the probability that one defector switches to cooperation $\Pr(B \rightarrow A)_G$ and the probability that one cooperator switches to defection $\Pr(A \rightarrow B)_G$.

$$\Pr(B \rightarrow A)_G = \left[1 + \exp(-\beta_G \cdot (d - a)) \right]^{-1} \quad (2.8a)$$

$$\Pr(A \rightarrow B)_G = \left[1 + \exp(-\beta_G \cdot (a - d)) \right]^{-1} \quad (2.8b)$$

Parameter β_G specifies the relation between the payoff difference of the members of the two different groups, and the probability of switching. Since payoffs are higher in cooperative groups than in defector groups ($a > d$), this process is expected to lead to the spread of cooperation between groups. Let j be the number of groups at the cooperative state, and let $m-j$ be the number of groups at the defector state. At state j , the probability of gaining (U_j^+) or losing (U_j^-) a cooperative group (by one defector taking over a cooperative group, or vice versa, by one cooperator taking over a defector group) can then be written as

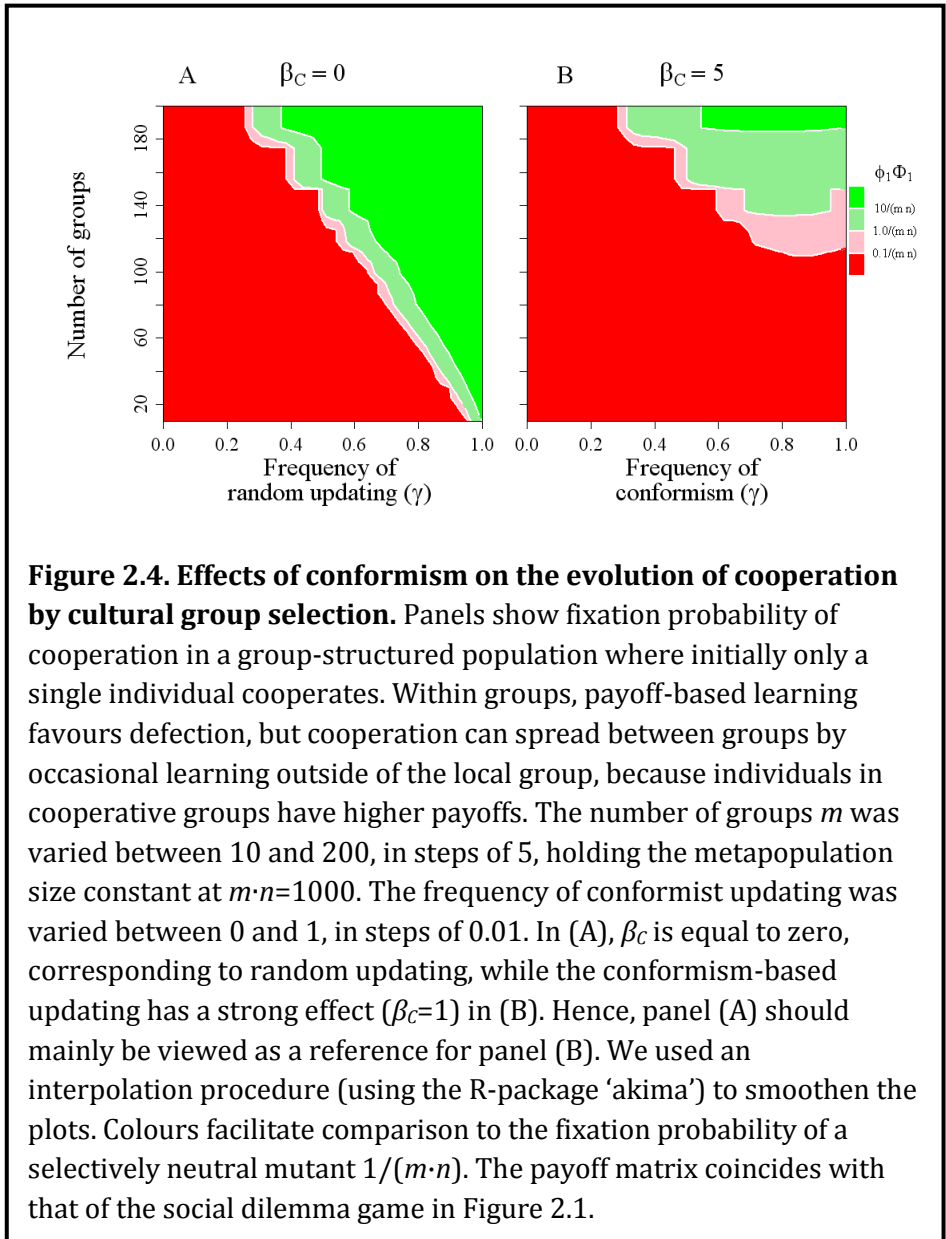
$$U_j^+ = \frac{j(m-j)}{m(m-1)} \cdot \varphi_1 \quad (2.9a)$$

$$U_j^- = \frac{j(m-j)}{m(m-1)} \cdot (1 - \varphi_{n-1}), \quad (2.9b)$$

where φ_1 refers to the fixation probability of a single cooperator in a group of defectors, and $1 - \varphi_{n-1}$ reflects the fixation probability of a single defector in a group of cooperators (which is complementary to the fixation probability of cooperation starting from state $n-1$). Using the same techniques as before, we can calculate the probability Φ_1 that eventually all groups have reached the cooperative state, given that we start out with 1 group of cooperators. The product $\varphi_1 \cdot \Phi_1$ then denotes the probability that cooperation reaches fixation in the metapopulation, given that we start out with one individual with the cooperative strategy.

Figure 4 gives an overview of the fixation probability of cooperation in case of a single cooperator in a population structured in m groups of fixed size n , under varying population structures and varying frequencies of conformism γ . When looking at Figure 2.4B only, one might conclude that conformism has a favourable effect on the cultural evolution of cooperation, since the fixation probability of cooperation tends to increase with γ . In our view, however, this conclusion would be misleading. The increasing scope for cooperation is not caused by conformism per se, but rather by the associated decrease in frequency of payoff-based learning. Decreased frequencies of payoff-based learning weaken selection against cooperation within the group. To assess the net effects of conformism, one therefore has to compare the results of Figure 2.4B with a benchmark that

takes this weakening of payoff-based learning into consideration. This benchmark is presented in Figure 2.4A where, with probability γ , individuals are imitating at random.



When the frequency of random switching within groups (as opposed to payoff-based switching) increases, the scope for cooperation increases (Figure 2.4A, going from the left to the right in the panel); within groups, the selection against cooperation is weakened, whereas selection between groups is kept constant (see Hauert and Imhof (2012) for a general analysis of how weakening within-group selection can affect the scope for cooperation in group-structured populations). Cooperation is favoured most, when the metapopulation is structured into many small groups (Figure 2.4A going from the bottom to the top of the panels). When groups are small, the probability that a single cooperative strategy reaches fixation – in spite of payoff-based learning disfavouring this strategy – is relatively large. Once such fixation has happened, cooperation can spread to other groups.

Comparing Figures 2.4A and 2.4B allows us to evaluate the net effects of conformism on the cultural evolution of cooperation by group selection. It is obvious that conformism hinders rather than favours the evolution of cooperation: in the whole parameter range, cooperation spreads more easily to fixation when conformism-based updating is random ($\beta_c=0$) than when it has a strong effect ($\beta_c=5$). Conformism hinders the evolution of cooperation because cooperation cannot gain a foothold in new groups: whenever a cooperator is introduced in a group of defectors, both conformist and payoff-based learning press an individual to switch back to defection.

Discussion

By means of a simple argument (where conformism was approximated by a coordination game), we have shown that the effect of conformism on cultural evolution strongly depends on the interaction context. In case of a coordination game, conformism merely affects the basins of attraction of the two pure-strategy equilibria; in case of a social dilemma, conformism can turn cooperation into a stable Nash equilibrium that coexists with an equilibrium corresponding to pure defection; and in case of an evasion game (such as a Hawk-Dove game or a Snowdrift game), conformism can destroy a polymorphic equilibrium and induce evolution to a pure-

strategy state. These conclusions were confirmed by a more mechanistic model for cultural evolution in a finite population. Including cultural group selection in this model strongly suggests that conformism tends to hinder, rather than promote the cultural evolution of cooperation by group selection.

Our approach using two-by-two matrix games with pure strategies is mathematically convenient, allowing to evaluate the success of strategies under a range of conditions in a fast and fairly straightforward manner. Also, this approach allows for deriving a diffusion approximation of the stochastic process, leading to a closed-form expression for fixation probabilities of a strategy under any mixture of payoff-based and conformist learning (see section 2 of the Appendix). Individual-based simulations in which switching between A and B is prone to errors – whose magnitudes are inversely related to β_p and β_c – lead to very similar outcomes in terms of fixation probabilities (not shown). This suggests that, despite the simplicity of our model, our findings are robust with respect to the way in which stochastic effects are introduced in the switching dynamics.

Our analysis is, however, restricted to the situation where each individual can only adopt a pure strategy. It remains unclear how conformism would influence cultural evolution when individuals are characterised by a mixed strategy, which specifies a probability distribution over the pure strategies. If such probabilistic tendencies could be transmitted between individuals by social learning, the dynamics of cultural evolution could be rather different from the scenario considered here (see (Gargiulo and Ramasco 2012) for an example). For instance, in an evasion game, all individuals could fixate on the same mixed strategy $0 < p < 1$, supporting a polymorphism where individuals make use of both pure strategies in a probabilistic fashion. In contrast to the findings from our analysis using two pure behaviours, conformism would be unlikely to destabilise such a behavioural polymorphism. Note, however, that the transmission of strategies by social learning depends on the degree to which individuals can evaluate the strategies of their peers. Whereas individuals might be able to evaluate the pure strategies of their peers (and possibly imitate

them accordingly), it is not obvious that more complex (mixed) strategies readily transmit between individuals.

Furthermore, we assume that individuals in the population all use the same social learning strategy. Decision making experiments show that individuals tend to vary considerably in their social learning strategies (*e.g.*, in the degree in which individuals learn based on payoffs; Molleman et al. in press; McElreath et al. 2005). Such individual variation in social learning strategies can affect the course of cultural evolution. To see this, consider a group in which some individuals typically learn based on payoffs and others learn based on conformism. This group may reach a stable behavioural polymorphism in an evasion game: when conformists all perform a common strategy, payoff-based learners can anticipate to that by adopting the strategy that is rare. Such an emerging differentiation, in which conformist learners perform one behaviour and payoff-based learners perform another behaviour, cannot be attained by groups that are homogeneous with respect to their social learning strategies. In our model, individuals use a mixture of conformist and payoff-based learning. As a consequence, a behavioural polymorphism is destabilised by a number of consecutive conformist learning events, potentially reducing the average payoffs of individuals in a population.

Human social learning comes in many different forms, and payoff-based and conformist learning only represent those forms that have received most attention in the social learning literature. Our analysis does not account for how other relevant forms of social learning – such as following a leader or a teacher – would affect the spread of behaviours within groups (see, *e.g.*, Lehmann and Feldman 2008 for how leadership can affect the cultural evolution of cooperation). Also, the mechanism that spreads behaviours between groups considered our model, is only one way that this group-level mechanism might work (Henrich 2004). Alternative scenarios in which groups of cooperators grow faster and split up when reaching a certain size (Traulsen and Nowak 2006) are likely to lead to different outcomes of cultural evolution. Such alternative forms of group-level selection can have different consequences for evolutionary dynamics, and can interact with within-group social learning in different ways. A more specific simulation study (Molleman et al. 2013), considering other population structures and different forms and social

learning and group selection, arrives at a similar conclusion: conformism can promote the cultural evolution of cooperation by group selection when groups can replace other groups, whereas cooperation cannot evolve when cooperators have to spread singly from group to group by a process of 'infection'.

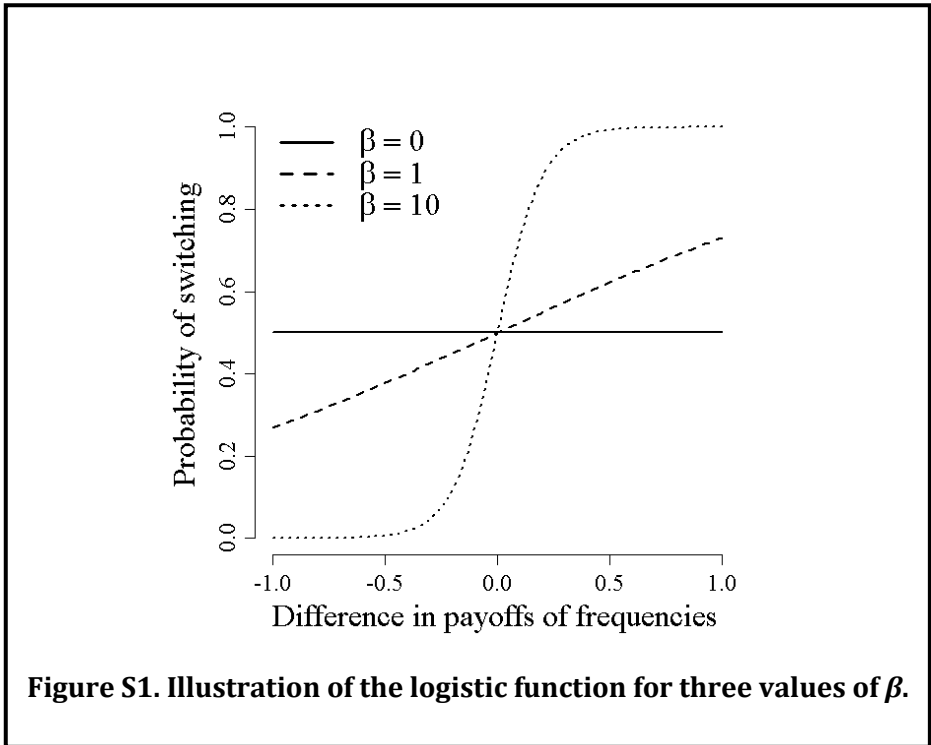
Our study leads to the conclusion that conformism has only a marginal effect in the context of coordination games, that it tends to erode the polymorphic equilibrium in evasion games, and that it does not favour cooperation in a social dilemma. In other words, our evaluation of the role of conformism is considerably less favourable than the opinion of other scholars of cultural evolution (Boyd and Richerson 1985; Henrich and Boyd 1998; Henrich and Boyd 2001; Henrich 2004; Andrés Guzmán et al. 2007). To put our conclusion into context, we would like to stress that our analysis did not consider potential intrinsic benefits of conformism. Two such benefits may be relevant in a variety of situations. First, conformism could homogenise groups with respect to a diversity of norms and habits, potentially increasing social cohesion and facilitating the establishment of trust, thus making it easier to resolve internal conflicts and to get cooperation off the ground. Second, conformism could be a beneficial strategy in environments with considerable spatial variation in payoffs and/or behavioural norms. In a situation like that, newcomers in a local group or environment could profit from imitating local habits, thus quickly adopting locally superior strategies or adapting to local behavioural equilibria. More sophisticated models accounting for spatial variation and the mechanisms underlying decision making in groups might therefore arrive at a more positive judgement of the role of conformism for the cultural evolution of social behaviour.

Appendix

In the first part of this Appendix, we provide a graphical illustration of the logistic functions used in our stochastic model to translate information on payoff or frequency differences into the probability of switching to another strategy. In the second part, we derive a diffusion approximation of our stochastic model that yields, by good approximation, analytical expressions for the fixation probability of a strategy in a finite population as a function of the strategy's initial frequency.

S2.1. Modelling the probability of switching to another behaviour as a function of payoff or frequency differences

The stochastic model presented in Section 2.4 of the main text specifies probabilities of switching as a logistic function of the difference in payoffs or frequencies of strategy A and B (equations 2.6a, b and 2.7a, b). Figure S2.1 below illustrates this relationship between the payoff or frequency differences and switching probabilities for three values of β . The function $[1 + \exp(-\beta \cdot x)]^{-1}$ specifies switching probabilities when two individuals with opposite strategies are paired, given the difference x in their payoffs or frequencies. Parameter β specifies the relationship between this difference and the probability of switching. When β is large, switching is strongly biased towards more frequent behaviour (or behaviours with higher payoffs). When β is small, switching is error-prone; a switch towards rare (or lower-payoff) behaviour will frequently occur by conformist (or payoff-based) learning.



S2.2. Diffusion approximation of our stochastic model

The stochastic model of the main text tracks the abundance i of behaviour A in a population of size n using transition probabilities between states i of the population. In each time step, this state i may change into $i+1$ or $i-1$, otherwise, the state remains unchanged. Here we use standard methods (Karlin and Taylor 1975; Otto and Day 2007, Chapters 14 and 15) to derive an analytical approximation of this discrete process by means of a diffusion model. The approximation transforms the discrete stochastic model in such a way, that the time and state space intervals become so small that the change of the system can be analysed as if it were continuous.

For switching probabilities, we take the definition of T_i^+ and T_i^- from the updating process described in section 3 of the Analysis and Results. In the diffusion model, it is convenient to work with the frequency $p = i/n$ of

behaviour A and the rescaled time $t' = t/n^2$. We have to derive the drift coefficient $\mu(p)$ and the diffusion coefficient $\sigma^2(p)$ for small time steps ($\Delta t' \rightarrow \infty$), which in the chosen parameterization corresponds to large values of n ($n \rightarrow \infty$).

Using the method in Otto and Day (2007, Box 15.2), it is straightforward to derive the first three moments of the expected frequency change per unit rescaled time:

$$m_1 = \frac{1}{\Delta t'} E(\Delta p \mid p = i/n) = n(T_i^+ - T_i^-) \quad (\text{S2.1})$$

$$m_2 = \frac{1}{\Delta t'} E(\Delta p^2 \mid p = i/n) = T_i^+ + T_i^- \quad (\text{S2.2})$$

$$m_3 = \frac{1}{\Delta t'} E(\Delta p^3 \mid p = i/n) = n^{-1}(T_i^+ - T_i^-) \quad (\text{S2.3})$$

In the limit $n \rightarrow \infty$, the third moment vanishes and the drift and diffusion terms are given by, respectively,

$$\mu(p) = \lim_{n \rightarrow \infty} m_1 = \frac{1}{2} p(1-p)[pz_1 + (1-p)z_2 - (1-2p)z_3] \quad (\text{S2.4a, b})$$

$$\sigma^2(p) = \lim_{n \rightarrow \infty} m_2 = p(1-p)$$

where the terms z_k are given by

$$\begin{aligned} z_1 &= (1-\gamma)(a-c)n\beta_p \\ z_2 &= (1-\gamma)(b-d)n\beta_p \\ z_3 &= \gamma n\beta_c \end{aligned} \quad (\text{S2.5, a, b, c})$$

The diffusion approximation is valid if the terms z_k remain finite for large values of n , *i.e.* when $n\beta_p$ and $n\beta_c$ remain bounded. In other words, β_p and β_c are of order $1/n$: payoff-based learning and frequency-based learning are both weak in large populations.

With the help of the drift and the diffusion coefficient, we can now express the fixation probability of A as a function of the initial frequency of A , $p_0 = k/n$. In general, this probability is given by the equation (Otto and Day, 2007, p.677):

$$\varphi_k = \frac{\int_0^{p_0} e^{-A(p)} dp}{\int_0^1 e^{-A(p)} dp}, \text{ with } A(p) = 2 \int \frac{\mu(p)}{\sigma^2(p)} dp. \quad (\text{S2.6})$$

Inserting $\mu(p)$ and $\sigma^2(p)$ calculated above reveals that φ_k can be expressed in terms of the error function (erf):

$$\varphi_k = \frac{\text{erf}(\rho_k/h) - \text{erf}(\rho_0/h)}{\text{erf}(\rho_n/h) - \text{erf}(\rho_0/h)} \quad (\text{S2.7})$$

with

$$\rho_k = (k/n)(2 \cdot (z_1 - z_2 + 2 \cdot z_3)) + z_2 - z_3 \quad (\text{S2.8a, b})$$

$$h = \sqrt{2 \cdot (z_1 - z_2 + 2 \cdot z_3)}$$

This approximation of fixation probabilities matches the fixation probabilities in the stochastic model quite well (see Figure 2.2 of the main text).

Chapter three

Cultural evolution of cooperation: the interplay between forms of social learning and group selection

Lucas Molleman
Andrés E. Quiñones
Franz J. Weissing

Evolution and Human Behavior 34, 342-349,
[doi:10.1016/j.evolhumbehav.2013.06.001](https://doi.org/10.1016/j.evolhumbehav.2013.06.001)

Abstract

The role of cultural group selection in the evolution of human cooperation is hotly debated. It has been argued that group selection is more effective in cultural evolution than in genetic evolution, because some forms of cultural transmission (conformism and/or the tendency to follow a leader) reduce intra-group variation while creating stable cultural variation between groups. This view is supported by some models, while other models lead to contrasting and sometimes opposite conclusions. A consensus view has not yet been achieved, partly because the modelling studies differ in their assumptions on the dynamics of cultural transmission and the mode of group selection. To clarify matters, we created an individual-based model allowing for a systematic comparison of how different social learning rules governing cultural transmission affect the evolution of cooperation in a group-structured population. We consider two modes of group selection (selection by group replacement or by group contagion) and systematically vary the frequency and impact of group-level processes. From our simulations we conclude that the outcome of cultural evolution strongly reflects the interplay of social learning rules and the mode of group selection. For example, conformism hampers or even prevents the evolution of cooperation if group selection acts via contagion; it may facilitate the evolution of cooperation if group selection acts via replacement. In contrast, leader-imitation promotes the evolution of cooperation under a broader range of conditions.

Introduction

The extension, degree and diversity of cooperation among unrelated individuals are key to the ecological success of humans. The term 'cooperation' refers to behaviours by which benefits arise from the interactions between individuals. Hence cooperative behaviour provides benefits at the group level. From the individual perspective, however, the evolutionary emergence and stability of cooperation is often puzzling. In particular, this holds for social dilemmas where performing a cooperative act is costly to the actor, and free-riding individuals can reap the benefits of cooperation without paying the costs.

In evolutionary biology, which is focused on genetic evolution, the evolutionary emergence and stability of cooperation is the subject of a considerable body of literature (Axelrod 1984; Lehmann and Keller 2006; Nowak 2006). Since the dawn of evolutionary theory, Darwin suggested that the evolution of cooperation might be explained by the differential performance of cooperative and non-cooperative groups in intergroup competition (Darwin 1859, 1871). Ever since then, this idea has been controversial (Maynard Smith 1964; Queller 1992; West et al. 2008; Leigh 2010). Evolutionary models demonstrate that selection between groups can indeed favour cooperation, but only under a limited range of demographic conditions (Maynard Smith 1964; Leigh 1983; Lehmann and Keller 2006; Lehmann et al. 2006; Traulsen and Nowak 2006). The problem is that within-group processes are typically faster than between-group processes. The rapid spread of individually favoured strategies (like refraining from cooperation) within groups erodes intergroup variation and, as a consequence, undermines the effectiveness of selection at the group level.

It has been argued that when social strategies are transmitted culturally rather than genetically, group selection can favour the evolution of cooperation under less restrictive conditions. The transmission of cultural traits is mediated by various forms of social learning, some of which play a key role in theories of cultural group selection. In particular, conformism, which is the individual tendency to acquire locally common strategies

(Boyd and Richerson 1985), can retard or prevent the spread of initially rare defective strategies (Henrich and Boyd 1998). By homogenizing behavioural strategies within groups, conformism changes the distribution of variation within and between groups, rendering cultural group selection a potentially efficient force promoting the evolution of cooperation (Boyd and Richerson 1985; Boyd and Richerson 2002; Bowles et al. 2003; Boyd et al. 2003; Henrich 2004; Guzmán et al. 2007; Scheuring 2009; Boyd et al. 2011). Recent models, however, cast doubt on the facilitating role of conformism in the evolution of cooperation (Lehmann et al. 2008; Lehmann and Feldman 2008; Peña et al. 2009). These models indicate that conformism can even hamper the evolution of cooperation because it hinders the spread of *any* new strategy in a group. Some modelling studies suggest that other mechanisms of social learning, such as the tendency to follow a group leader are more efficient in promoting cooperation via cultural group selection (Cavalli-Sforza and Feldman 1981; Lehmann et al. 2008; Lehmann and Feldman 2008). The contrasting conclusions of different modelling studies by different schools of thought have led to an on-going debate on the role of cultural group selection. Part of the debate centres on the range of parameter values considered reasonable by different authors. Perhaps more importantly, comparison across models is hampered by the fact that the models differ in their basic assumptions on social learning and group selection (Boyd et al. 2011).

In the theory of genetic evolution, there has been considerable confusion and debate around the meaning of 'group selection'. Some authors even argue that this term should be abandoned altogether (West et al. 2007), since group selection is a special case of a more general hierarchical theory of selection (Frank, 1986). Yet the concept of group selection can be useful in the context of a group-structured population, since group traits may be an emergent property of the interaction of group members that cannot easily be reduced to individual-level traits (Okasha 2006; Simon et al. 2012). When talking about group selection, it is crucial to distinguish between two main modes, each representing different mechanisms that can have different effects on the outcome of evolution (West et al. 2007). First, group selection can be driven by the differential extinction and colonization of patches, that is, the replacement of less

successful groups by more successful ones (Maynard Smith 1964; Leigh 1983). This mode of group selection played a prominent role in the earliest discussions on group selection and is currently been referred to as 'old' group selection (West et al. 2007) or multilevel selection 2 (Okasha 2006). We prefer to use a more descriptive term and will henceforth refer to 'replacement group selection' when group selection is driven by the replacement of less successful groups by more successful ones. Second, group selection can be driven by the differential production of individuals migrating to and settling in other groups, that is, the 'contagion' of groups by individuals derived from other groups (Wilson 1975; Rogers 1990). This mode of group selection has been coined 'trait-group' selection (Wilson 1975), 'new' group selection (West et al. 2007), or multilevel selection 1 (Okasha 2006). As a more descriptive term, we propose to call this mode 'contagion group selection'. In the biological literature, the distinction between group selection by replacement and group selection by contagion, or the lack of it, has produced extensive discussion (Okasha 2004, 2006; West et al. 2007, 2008; Wilson 2008) and confusion (Wilson and Wilson 2007). Despite this, there have been very few attempts to compare the requirements for each of these processes to work (Lehmann et al. 2006; García and Van den Bergh 2011).

In cultural evolution, the distinction between replacement and contagion group selection is as relevant as in biological evolution (Henrich 2004). Replacement group selection corresponds to the cultural take-over of whole groups by other more successful groups (Boyd et al. 2003; Traulsen and Nowak 2006; Guzmán et al. 2007). This may happen as a result of intergroup conflicts, where the winning group imposes their 'culture' upon subdued groups (Boyd et al. 2003; Traulsen and Nowak 2006; Guzmán et al. 2007). It may proceed in a more indirect way, if less successful groups tend to disband and go extinct, while well-performing groups bud off subgroups recolonising empty patches (Soltis et al. 1995). Contagion group selection is mediated by the more gradual migration of cultural traits from one group to another. This may, for example, happen if the cultural traits observed in well-performing groups are preferentially imitated by the individuals of other groups, leading to the gradual introgression of group-beneficial strategies into less successful groups (Henrich and Gil-White 2001; Boyd and Richerson 2002; Lehmann et al.

2008). It is to be expected that, as in genetic evolution, the two modes of group selection have contrasting effects on the course and outcome of cultural evolution. Yet, systematic studies on these effects are currently lacking.

Before continuing it is important to clarify the meaning of the term 'selection' in the context of cultural evolution. We use a definition that is analogous to the usage of natural selection in genetic evolution, but somewhat more restricted than the definitions often given in the literature on cultural evolution (*e.g.*, Cavalli-Sforza and Feldman 1981; Boyd and Richerson 1985). In genetic evolution, natural selection refers to the preferential transmission of alleles that enhance the 'fitness' of individuals, groups or other entities, that is, the ability of these entities to survive and reproduce. Darwinian fitness is often measured in payoffs that are acquired in interactions with others. Natural selection is an important driver of evolutionary change, but there are many other processes (including genetic drift and mutation pressure) leading to a change in allele frequencies. These processes differ from natural selection in that the evolutionary success of an allele is not related to the effect this allele has on the performance of the individuals, groups or other entities harbouring this gene. Similarly, in cultural evolution, various processes lead to changes in frequencies of cultural variants. In some of these processes, the evolutionary success of a cultural variant is related to the 'performance' of individuals, groups or other entities harbouring this variant. Examples include the preferential imitation of high-payoff individuals or the higher rate of cultural take-over from groups with superior organisation or technology. However, as in genetic evolution, the cultural transmission of a trait is not necessarily linked to the 'performance' of this trait, or to any of its inherent properties. Examples of forms of social learning that lead to cultural changes independent of payoffs include conformism-based learning (imitating traits that are locally most frequent), and status- or reputation-based learning (*e.g.*, following a leader irrespective of intrinsic qualities of his/her cultural traits). To maintain consistency with genetic evolution, we interpret cultural change that is driven by performance- or payoff-based social learning as analogous to natural selection, while cultural change driven by other forms of transmission (such as conformism) does not fall into this category.

In this paper, we study three forms of social learning: payoff-based learning, conformism and leader imitation. Firstly, humans preferentially copy strategies from successful individuals (Henrich and Gil-White 2001; Lehmann et al. 2008; Kendal et al. 2009). In the context of a social dilemma, this social learning rule is expected to decrease rates of cooperation within groups since free-riding leads to higher payoffs. Secondly, under conformist learning individuals tend to adopt locally common strategies, thereby further increasing the frequency of those strategies in their groups. Thirdly, individuals may be inclined to follow a leader or a teacher in their group, so that the strategy of one influential individual tends to spread locally (Cavalli-Sforza and Feldman 1981). Experimental results suggest that more than one of the abovementioned social learning rules can apply at the same time (Efferson et al. 2008; McElreath et al. 2008).

To clarify matters and to help resolving the disagreements in the literature on the role of cultural group selection, we developed an individual-based model of cultural evolution in a group structured population. The model allows to consider various mixtures of social learning rules (payoff-based learning, conformism, leader imitation) in the context of two contrasting modes of cultural group selection (contagion versus replacement). For all combinations of settings, we systematically varied the strength of individual and group selection with the aim of elucidating whether, and under which conditions, cultural group selection favours the evolution of cooperation in a social dilemma.

The model

3.1. Overview

We simulate a metapopulation structured in m groups of n individuals each. Individuals face a social dilemma, modelled here as a public goods game. At each point in time, each has a pure strategy, 'cooperate' or 'defect'. Cooperators contribute to the public good of their group at a cost. Defectors contribute nothing. All contributions to the public good yield a benefit that is distributed equally among all group members, irrespective of whether they have contributed to the public good or not. In our

simulations of cultural evolution, payoffs acquired in the public goods game influence the probability that a strategy spreads due to payoff-based learning. In each time step of the simulation, one individual ('focal') is randomly sampled from the metapopulation to change its strategy by imitating one or more other individuals. Imitation can be based on payoffs, conformism to the majority, or following a leader. After each iteration, the payoff that individuals obtain from the game is updated. This stepwise procedure leads to an evolutionary dynamics in which the population can lose or gain one cooperator per iteration. We iterate this event-based process until cooperation has either been lost in the metapopulation or spread to fixation. Imitation of strategies can either occur within the focal individual's group or in the context of a different group. The probability of occurrence of within and between group imitation events is given by the parameter γ that corresponds to the probability of a group level event. Within-group imitation occurs either based on individual payoffs from the public goods game (with probability α) or based on another imitation rule (probability $1 - \alpha$). As alternatives to payoff-based learning, we consider conformism, leader-following and – as a standard of comparison – the imitation of a randomly chosen group member. As motivated in the introduction, we interpret changes in the frequency of strategies due to payoff differences as 'selection'. Therefore, the parameter α that specifies the frequency of payoff-based imitation within groups represents the strength of individual selection. Between-group imitation occurs either randomly (with probability $1 - \beta$), or based on benefits of the public goods in the respective groups (probability β). The parameter β specifies the frequency of payoff-based imitation at the group level, and therefore the strength of group selection. Events involving different groups are implemented in two ways, reflecting two modes of group selection (contagion versus replacement). Figure 3.1 gives a graphical representation of the possible events that occur in a single iteration of the simulation. Details of the implementation of the imitation rules and modes of group selection are given below. First we examine scenarios where the metapopulation initially consists of one fully cooperative group and defectors in all the other groups. This is an assumption that has been adopted in several models of cultural group selection (Henrich and Boyd 2001; Boyd et al. 2003; Guzmán et al. 2007) and is based on the idea that cooperation gained a foothold and has spread to fixation in one group due

to stochastic processes. Second, we contrast these scenarios with simulations that start with a whole metapopulation of defectors, allowing cooperation to arise spontaneously with a fixed probability (corresponding to a mutation in genetic evolution).

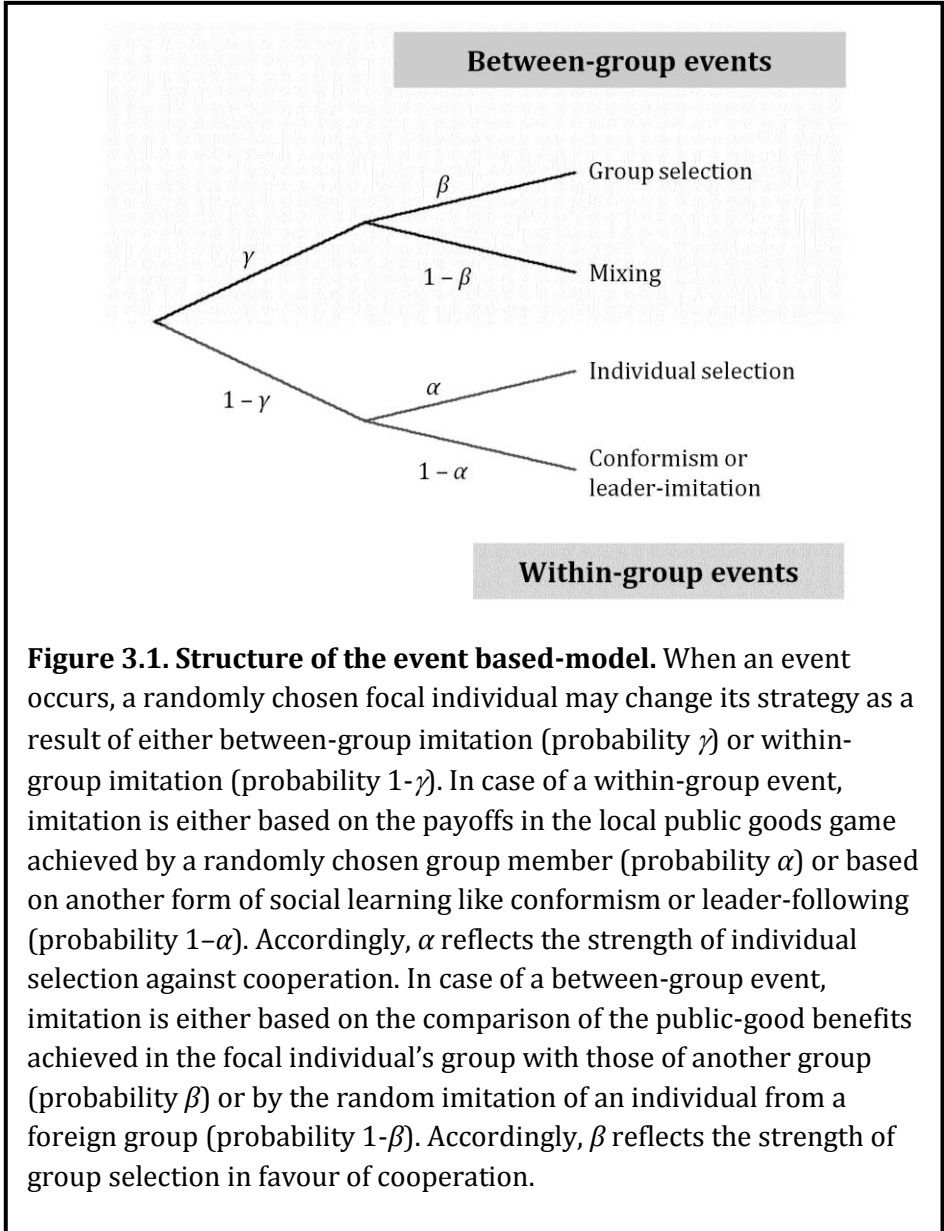


Figure 3.1. Structure of the event based-model. When an event occurs, a randomly chosen focal individual may change its strategy as a result of either between-group imitation (probability γ) or within-group imitation (probability $1-\gamma$). In case of a within-group event, imitation is either based on the payoffs in the local public goods game achieved by a randomly chosen group member (probability α) or based on another form of social learning like conformism or leader-following (probability $1-\alpha$). Accordingly, α reflects the strength of individual selection against cooperation. In case of a between-group event, imitation is either based on the comparison of the public-good benefits achieved in the focal individual's group with those of another group (probability β) or by the random imitation of an individual from a foreign group (probability $1-\beta$). Accordingly, β reflects the strength of group selection in favour of cooperation.

3.2. The public goods game

The social dilemma is modelled as a public goods game with discrete strategies. Cooperators contribute c units of their resources to a public good. Defectors do not contribute, avoiding the cost of cooperation. The total of investments is multiplied by $a > 1$, yielding the overall benefit $b = n_c \cdot a \cdot c$ to a group containing n_c co-operators. This benefit is distributed equally among the n members of the group, irrespective of their contribution. Accordingly, payoffs acquired by cooperators and defectors are $P_c = b/n - c$ and $P_d = b/n$, respectively. Since the payoff of a defector is always higher than that of a co-operator, within-group selection favours defection over cooperation.

3.3. Within-group events

The focal individual can imitate a member of its own group with probability $1 - \gamma$. We consider four social learning rules determining this within-group imitation. With probability α , social learning is based on the payoffs acquired in the public goods game. With probability $1 - \alpha$, payoff-based imitation is complemented by either random imitation, conformism or leader-imitation. These four learning rules were implemented as follows:

- Payoff-based imitation: the payoff of the focal individual P_f is compared with the payoff P_m of a randomly chosen group member (the ‘cultural model’). The focal individual switches its strategy to that of the cultural model when the perceived payoff difference $P_m - P_f + \varepsilon$ is positive, where the ‘noise’ term ε is drawn from a standard normal distribution.
- Random imitation: the focal individual switches its strategy to that of a randomly chosen group member. Random imitation corresponds to genetic drift in genetic models of evolution. It is included in our model as a benchmark to assess the effects of a decrease in the strength of individual selection α on the evolution of cooperation in the absence of conformism and leader-imitation.

- Conformism: the focal individual samples three randomly chosen models from its group and imitates the majority in that sample (Boyd and Richerson 1985). With only two strategies in our model, taking three models guarantees a majority in a sample of minimal size. This small sample size reflects limited availability of information (*e.g.*, due to constraints on the time that individuals can spend sampling the behaviour of other group members), or limited cognitive capacities. Alternative implementations of conformism, including larger sample sizes, are discussed below.
- Leader-imitation: the focal individual imitates the ‘leader’ of its group. Each group has a leader, who is randomly chosen at the beginning of the simulation. The identity of the leader remains unchanged throughout a simulation run. Leaders may change their strategy by imitation, like any other member of their group by payoff-based imitation or between-group events. Leader-imitation homogenises strategies within groups, irrespective of their payoff or frequency.

3.4. Between-group events

Imitation based on between-group comparisons occurs with frequency γ . With probability $1 - \beta$, the focal individual imitates a randomly chosen individual from any group in the population. This process mimics the random migration of strategies or mixing of ideas between groups. With probability β , imitation is based on a measure of group performance, reflected in the acquired payoffs. In this case, the focal individual compares the per-capita benefit b_f/n from the public goods in its own group to the per-capita benefits b_m/n in a randomly chosen ‘model group’. The perceived benefit difference is $(b_m - b_f)/n + \varepsilon$, where the ‘noise’ term ε is again drawn from a standard normal distribution. We consider the following modes of group selection:

- Contagion-based group selection: if the perceived benefit difference is positive, the focal individual imitates a randomly chosen individual from the model group. The group of the focal individual is ‘infected’ with the strategy of the model group.

- Replacement-based group selection: if the perceived benefit difference is positive, with probability $1/n$ the strategies of all n members of the group of the focal individual is changed to the strategies in the model group. The group of the focal individual is replaced by the model group, so that the group of the focal is an exact copy of the model group. With probability $1 - 1/n$, no change does occur. We introduced the probability $1/n$ to compensate for the fact n group members change their strategies in case of group replacement. This allows a more direct comparison of contagion-based and replacement-based group selection, as the same number of individuals is expected to change in every time step.

3.5. Initial conditions

In the first set of simulations, we initialise our population with one group composed of only cooperators and the other $m-1$ groups composed of only defectors. This choice of the initial state rests on the assumption that one group has shifted to a cooperative state due to stochastic effects (Boyd et al., 2003, 2011). The initial presence of strategic variation allows us to study the effect of transmission and selection without ‘mutation’. In the second set of simulations, we start with m groups of defectors (and no cooperators) to see if a stochastic group shift is likely to occur. To this end, we allow individuals to spontaneously change their strategy with probability μ . While we consider this second set of simulations - including spontaneous changes - more realistic, we include both sets to compare the effects of common settings used in the literature.

In all simulations, we track the evolution of cooperation by iterating the model for a maximum of 500 000 time steps (events). Simulation runs finished when either cooperation or defection had reached fixation in the population. Parameters (α, β, γ) were varied between 0 and 1 with steps of 0.1, and with steps of 0.01 for values under 0.1. Based on preliminary simulations runs, we chose to run 30 replicates for each parameter combination to assure that the results were robust to stochastic processes. In the results reported in the main text, we consider a metapopulation of $m=50$ groups consisting of $n=20$ individuals, yielding a

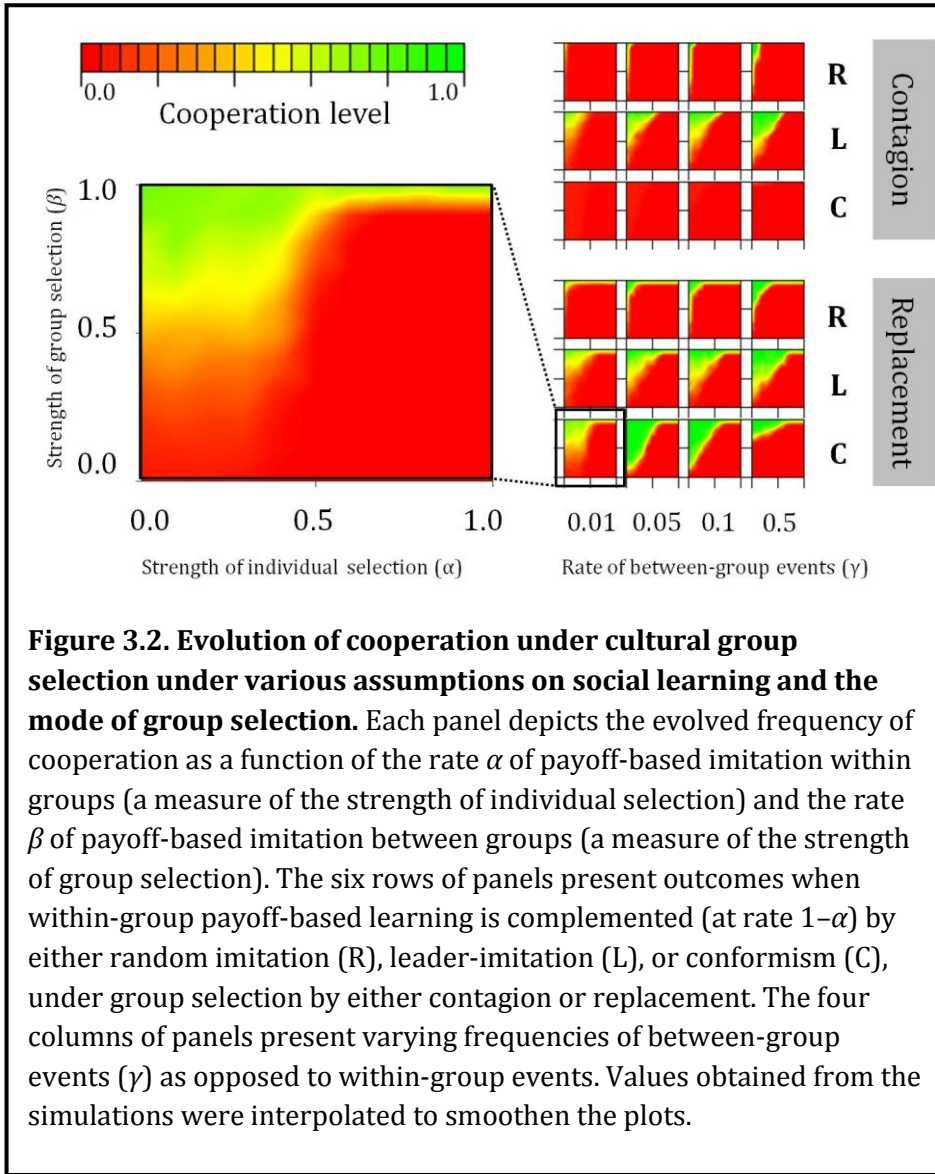
total population size of 1000 individuals. All figures shown below are based on the payoff parameters $a=3$ and $c=1$.

Results

For each parameter combination, we calculated the frequency of cooperation in the metapopulation at the final time step averaged across the 30 replicate simulations. Figure 3.2 gives an overview of the simulations that were initialised with one cooperative group. Each panel summarizes the outcome of cultural evolution as a function of the parameters α (strength of individual selection) and β (strength of group selection). The 24 small panels to the right correspond to the scenarios obtained by combining two modes of group selection (contagion-based versus replacement-based) with three alternatives to payoff-based imitation (random imitation R, leader-imitation L, conformism C) and four levels of γ (corresponding to the relative frequency of between-group events).

Selection strengths

As exemplified by the enlarged panel in Figure 3.2, the evolved frequency of cooperation decreased with the strength of individual selection (from left to right), while it increased with the strength of group selection (from bottom to top). In all scenarios considered, cooperation disappeared in case of weak group selection ($\beta \approx 0$) while it spread to fixation in case of strong group and weak individual selection ($\beta \approx 1, \alpha \approx 0$). This makes sense, since payoff-based imitation disfavors cooperation in case of within-group events, while it favors cooperation in case of between-group events. Strong individual selection ($\alpha \approx 1$; right-hand side of each panel) almost invariably lead to the disappearance of cooperation; only in case of strong ($\beta \approx 1$) replacement-based group selection it remained in the population and even spread to fixation.



The scenarios where payoff-based imitation is combined with random imitation (panel rows R) exemplify how a decrease in the strength of individual selection on its own may lead to the establishment of cooperation. The corresponding panels serve as a benchmark for the cases where payoff-based imitation is combined with leader-imitation (panel rows L) or with conformism (panel rows C). A comparison across panel

rows reveals that the effects of leader-imitation and conformism, while clearly present, only marginally expand the parameter range for the establishment of cooperation. In case of contagion-based group selection, conformism even leads to a reduction in the parameter range allowing for the evolution of cooperation. In that sense, conformism can hamper the evolution of cooperation by cultural group selection.

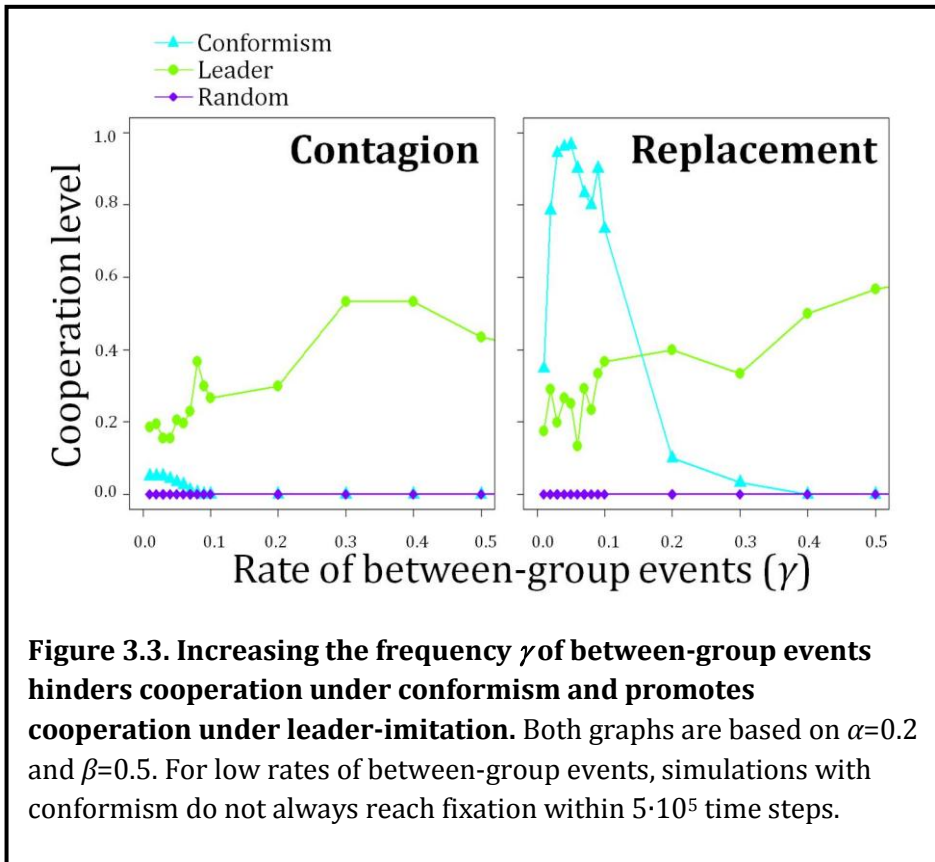
Group selection by contagion

When group selection is based on contagion, only leader-imitation favours the evolution of cooperation, whereas conformism has a detrimental effect. It is easy to understand why conformism hampers the spread of cooperation: even if some individuals get 'infected' by a cooperative strategy (by imitating an individual from a group with high public-goods benefits), this strategy cannot get a foothold in the local group since within-group payoff comparisons act against cooperation while within-group conformism tends to weed out cooperation as long as it is rare. In contrast, leader-imitation can promote the evolution of cooperation provided that group selection is strong enough to overcome the effects of individual selection.

Group selection by replacement

Under group selection by replacement, both conformism and leader-imitation widen the scope for cooperation (Figure 3.2). In contrast to the contagion scenario, conformism can promote cooperation here. This is even the case when group selection is relatively weak (low values of β), particularly when within-group events are based on considerable frequencies of conformism (say, $\alpha < 0.4$). Note that conformism favours cooperation most when rates of between-group events (γ) are relatively low. This effect is also exemplified by Figure 3.3. When between-group events are rare, conformism can prevent defectors from invading in cooperative groups and facilitate the evolution of cooperation under group replacement (Figure 3.3, right panel). This effect breaks down if between-group events are too frequent. All else being equal, the impact of group selection increases with γ . However, beyond a certain level of γ , within-group dynamics is too slow for conformism to prevent the spread of defectors that are introduced by random mixing between groups. When

conformism is infrequent, variation between groups breaks down, weakening the potential of group selection to promote cooperation. While the effect of conformism on the evolution of cooperation is strongly dependent on the mode of group selection, leader-imitation has similar effects under contagion- and replacement-based group selection (Figure 3.2, panel rows L; Figure 3.3). In contrast to conformism, leader-imitation promotes cooperation more strongly under increasing rates of between-group events (Figure 3.2, panel rows L; Figure 3.3). The arrival of migrants from other groups does not diminish the homogenising power of leader-imitation.



Spontaneous emergence of cooperation

All results reported so far were obtained from simulations initialised with one group of cooperators. We relaxed this assumption by initialising the population with only defectors and allowed individuals to spontaneously switch their strategy by innovation. In this scenario, conformism does not promote the evolution of cooperation. Even under group selection by replacement, conformism impedes the evolution of cooperation, preventing the spread of any new strategy that arises by innovation (Figure 3.4). For leader-imitation, the initial conditions do not influence the outcome of cultural evolution qualitatively.

Robustness of results

In addition to the simulations reported above, we ran numerous other simulations. Changing the payoff parameters in the public goods game, the distribution of individuals over groups (*e.g.*, $m=10$, $n=100$; $m=40$, $n=25$) and the total population size (with $n=20$ and $m=500$) only had a small quantitative effect but did not change our conclusions. In contrast to the findings reported by Boyd and colleagues (2011), the outcome of cultural group selection was also not affected when we considered a spatially explicit model with stepping-stone migration on a torus (Boyd and Richerson 2002). More localized migration only slows down the speed of evolution, but does not change the outcome in a qualitative way.

In our implementation of conformist learning, an individual samples only three models in its group (Boyd and Richerson 1985). Increasing the size of this sample strengthens conformist effects, leading to more efficient within-group homogenization of behaviour (simulation data not shown). Moreover, alternative approaches to modelling the effects conformism on the cultural evolution of social behaviour yield results in line with the findings reported here (Molleman et al. 2013), suggesting that our simulation results are robust to changes in the way conformism is implemented.

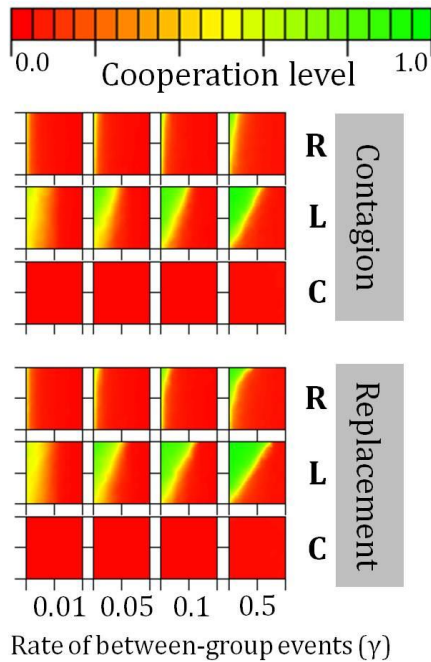


Figure 3.4. Outcome of cultural group selection when cooperation was initially absent and could only arise by spontaneous 'mutation' (innovation). Panels as in Figure 3.2. Per event, the focal individual spontaneously changes its strategy with probability $\mu=0.01$. Values obtained from the simulations were interpolated to smoothen the plots.

Discussion

The last years have seen a debate about the role of social learning in the evolution of cooperation by cultural group selection. On the one hand, it has been argued that a social learning rule like conformism favours the evolution of cooperation (Boyd and Richerson 2002; Boyd et al. 2003; Choi and Bowles 2007; Boyd et al. 2011). This argument is based on the idea that conformism reduces behavioural variation within groups, thereby increasing the effects of variation between groups (Henrich and Boyd 1998). As a result, cultural group selection could potentially be an

important factor facilitating the evolution of cooperation. On the other hand, recent theoretical work has challenged this argument, showing by means of example models that conformism does often not favour the evolution of cooperation (Lehmann et al. 2008; Lehmann and Feldman 2008). Our systematic comparison indicates that the contrasting conclusions reflect the specific combination of social learning rules and mode of group selection assumed in the models underlying these conclusions. Under group selection by contagion, individuals from cooperative groups are imitated more and defector groups get 'infected' by cooperative strategies. Conformism hinders the evolution of cooperation in this scenario: uncommon behaviours are strongly selected against and rare cooperators infecting defector groups are disfavoured by both payoff-based imitation and conformism. When cooperative groups can replace less cooperative groups, conformism can promote the evolution of cooperation, provided that initially one group in the population consists of cooperators (Boyd et al. 2003).

Our results contrast with previous models in terms of the facilitating effects of conformism on the evolution of human cooperation (Boyd and Richerson 1985; Henrich and Boyd 1998; Henrich and Boyd 2001; Boyd et al. 2011). However, one could think of scenarios in which conformism could have a positive influence on cooperation in a more indirect way. For instance, conformism might homogenize groups with respect to various norms and habits and thereby increase social cohesion. As a consequence, trust may build up among group members more easily, potentially facilitating cooperation. Secondly, conformism may have positive effects on cooperation when payoffs of behaviour vary spatially. Conformist learning can help newcomers to adopt locally beneficial strategies, allowing them to coordinate with resident individuals, and adapt to local equilibria (Boyd and Richerson 1985). Therefore, it is possible that under these different conditions conformism plays a more prominent role in the evolution of social behaviours. Further theoretical work is needed to clarify these issues.

Earlier studies (Lehmann and Feldman 2008; Lehmann et al. 2008) arrived at the conclusion that leader-imitation is more efficient than conformism in establishing cooperation through cultural group selection. Our results are in line with this conclusion. In fact, imitating a group

leader can promote the evolution of cooperation irrespective of the mode of group selection. However, the way that leadership is implemented in these models as well as ours, is a simplification that is certainly quite unrealistic. Leaders are most likely not chosen at random, but they emerge in the interaction between individuals (*e.g.*, Johnstone and Manica 2011; Weissing 2011). Becoming a leader or a follower and accepting the leadership of somebody else will often be the outcome of a game with strategies and payoffs that are quite unrelated to the public goods game considered here. Leadership and followership can also be institutionalized, again associated with costs and benefits that are not necessarily congruent with those of the underlying public goods game. Accordingly, leaders and followers will often be motivated by incentives that are not conceptualized in our simple model. It is easy to imagine situations where leader-imitation is an even more potent force in bringing about cooperation than in our model (*e.g.*, if institutions reward leaders on the basis of group benefits). However, in other situations leader-imitation might actually hamper the evolution of cooperation (*e.g.*, if the leader has to be paid from public-good benefits). Clearly more refined models of leadership are required to really judge the role of leader-imitation for the evolution of cooperation.

In total, we conducted more than $4 \cdot 10^9$ simulations, each simulation running for up to $5 \cdot 10^5$ time steps. Yet, even an extensive study like this can only address few potential interactions between social learning rules and modes of cultural group selection. For example, we assessed the effects of only two imitation rules (conformism and leader-imitation). There may be many other rules potentially reducing variation within groups (Laland 2004; Rendell et al. 2011), and the way this reduction comes about could interact with the mode of group selection in unexpected ways (witnessed by the effects of conformism reported here). Also, different modes of group selection are by no means mutually exclusive, and may act simultaneously and interactively (Okasha 2006). In our model, payoff-based imitation is error prone, but we did not systematically investigate the implications of various error rates and degrees of noise. For simplicity, we assumed that the errors associated with the measurement of payoff differences are equally large when imitation occurs within or between groups. In many situations it is more

plausible to assume that the success of individuals from other groups is harder to assess than the payoffs of group members.

Models studying the evolution of social learning mainly focus on contexts in which outcomes of behaviour are independent of the behaviour of others (Boyd and Richerson 1985). The same is true for most (McElreath et al. 2005; Efferson et al. 2007; McElreath et al. 2008; Efferson et al. 2008) but not all (Traulsen et al. 2010) experimental work investigating how humans use social information to determine their behavioural strategies. Our understanding of the cultural evolution of cooperation would benefit from theory that predicts which forms of social learning are adaptive in contexts where outcomes of behaviour also depend on the strategic choices of others. Confronting humans with such contexts under controlled laboratory conditions could then test if humans indeed use these social learning rules in decision making in cooperative interactions. Also, studying cooperation should not restrict itself to social dilemmas. Human cooperation comes in many different forms and flavours; the public goods game modelled here represents only one specific context in which humans cooperate. Exploring the effects of various social learning rules on the cultural evolution of social behaviour in other games, such as coordination games and evasion games, presents an interesting venue of research for the future.

Chapter four

Consistent individual differences in human social learning strategies

Lucas Molleman
Pieter van den Berg
Franz J. Weissing

Nature Communications, in press

Abstract

Social learning has been key to human evolutionary success. The cultural transmission of information between individuals has enabled human groups to adapt to a vast array of habitats. Using the ability to learn from their peers, humans have built up extensive cultural repertoires tailored to the environmental and social conditions that they faced. However, it is unclear which social learning strategies people use, especially in social contexts where their payoffs depend on the behaviour of others. Here we show experimentally that individuals differ in their social learning strategies and that they tend to employ the same learning strategy irrespective of the interaction context. Payoff-based learners focus on their peers' success, while decision-based learners disregard payoffs and exclusively focus on their peers' past behaviour. These individual differences may be of considerable importance for cultural evolution. By means of a simple model we demonstrate that groups harbouring individuals with different learning strategies may be faster in adopting technological innovations, and can be more efficient through successful role differentiation. Our study highlights the importance of individual variation for human interactions and sheds new light on the dynamics of cultural evolution.

Introduction

Human success in colonising nearly all terrestrial habitats of our planet was facilitated by our social nature (Boyd et al. 2011; Whiten and Erdal 2012). Humans excel in collective action, are able to cooperate in social dilemmas and to employ high degrees of group coordination and cooperation to solve adaptive problems. The transmission of information between individuals through social learning was a key factor for the spread of humans across the whole planet (Boyd et al. 2011). Such cultural transmission has led to the accumulation of huge amounts of adaptive information in very little time. It has been argued that social learning leads to a parallel inheritance system giving rise to a process that is in many ways analogous to genetic evolution (Dawkins 1976; Cavalli-Sforza and Feldman 1981; Boyd and Richerson 1985). The dynamics and outcome of this process of cultural evolution is to a large extent determined by the rules that govern the transmission of cultural information between individuals; *i.e.*, the social learning mechanisms that people use. Several social learning strategies have been described that may each be adaptive under particular circumstances (Laland 2004; Galef Jr and Laland 2005; Valone 2007; Kendal et al. 2009; Rendell et al. 2010; Rendell et al. 2011). A social learning strategy can specify under which circumstances to pay attention to social information; which individuals to choose as target for collecting information; which type of information to gather from these individuals; and how to use this information as guidance for future behaviour (Laland 2004; Hoppitt and Laland 2013). Prominent examples of social learning strategies are ‘conformist learning’ (to imitate a local majority; Boyd and Richerson 1985; Rogers 1988; Henrich and Boyd 1998; Richerson and Boyd 2004; Kendal et al. 2009), ‘payoff-based learning’ (to imitate the behaviour of peers who achieved high payoffs in the recent past; Boyd and Richerson 1985; Schlag 1998; Henrich and Gil-White 2001; Lehmann et al. 2008; Lehmann and Feldman 2008), or ‘prestige-based learning’ (to imitate or follow the advice of high-status individuals such as leaders or teachers; Henrich and Gil-White 2001; Lehmann and Feldman 2008; Lehmann et al. 2008).

Empirical evidence from the field suggests that these learning strategies do indeed affect the spread of behaviour in human groups (Henrich and Henrich 2010; Henrich and Broesch 2011). The results of such investigations are corroborated by evidence from controlled laboratory studies, indicating that people attend to both the frequencies of their peers' behaviour, as well as to the payoffs associated with it (McElreath et al. 2005; Efferson et al. 2007; McElreath et al. 2008; Efferson et al. 2008; Mesoudi 2008; Traulsen et al. 2010; Mesoudi 2011). In addition, there are indications that the extent to which people resort to social information depends on factors like task difficulty, confidence in their own behaviour (Morgan et al. 2012) and environmental variability (Toelch et al. 2009). Remarkably, experimental results suggest that individuals differ in the extent to which they rely on social information when making decisions under uncertainty (McElreath et al. 2005; Mesoudi 2011), and vary in their tendency to conform to behaviour of the majority (Efferson et al. 2008).

As yet, there is no strong evidence that individuals differ with respect to the type of social information that they rely on in making decisions. This is not surprising, since experimental studies of human social information use have predominantly focused on the spread of technological innovations that are mainly relevant in non-social contexts. For instance, when attempting to acquire technological know-how, individuals almost exclusively use payoff-based learning (by copying behaviour of successful individuals) when explicitly presented with a choice between multiple social learning strategies (Mesoudi 2011). Here we focus on a different type of problem: finding the adequate behaviour in a social setting. Such social 'know-how' differs from technological know-how in that the payoffs of the various alternatives do not only depend on the external environment, but also on the behaviour of others in the population. When having to choose among alternatives in a non-social setting, the information obtained from others is mainly useful for obtaining a more accurate estimate of the payoffs of these alternatives in a noisy world. In a social setting, these payoffs reflect the behaviour of others. As a consequence, social learning gets a new dimension, namely acquiring information on the distribution of behavioural tendencies in the population. The importance of such frequency information strongly

depends on the type of social interaction. For this reason, we here study social learning in social interactions with a different signature. This way we get a more general impression of whether and how social learning differs between technological and social contexts. Moreover, the inclusion of different types of social context allows us to investigate whether individuals are consistent across those contexts. Do social learning strategies reflect general behavioural tendencies, or do individuals flexibly adjust their use of social information to the type of social interaction at hand?

To obtain a better understanding of human social learning in both social and non-social settings, we designed two experiments. Subjects were grouped and repeatedly made decisions that determined the amount of money earned during the experiment. Before making a decision, they were allowed to request information about the behaviours and payoffs of their peers. We evaluate social learning strategies by tracking the types of information requests that subjects made.

In eight independent sessions of our first experiment (henceforth called the 'primary experiment'), we confronted sixteen subjects ($n=128$) with four settings ('contexts'). In each context, the subjects had to make binary choices between options *A* and *B*, but the contexts differed in the way payoffs depended on the subject's own decision and the decisions of others. We based our design on the experimental paradigm of McElreath et al (2005), and extended this setup to include interactions in the three strategically different classes of games with pure strategies (Molleman et al. 2013a), each of which has been extensively researched in the literature on the evolution of cooperation (Gintis 2000). Each subject sequentially encountered *i*) a best choice (BC) situation where one options (*B*) yields on average a higher payoff; *ii*) a social dilemma (SD) where one option (*A*) increases the payoffs of all group members but in each situation yields a lower payoff than the other option (*B*) for each individual employing it; *iii*) a coordination game (CO), where the payoff of each option (*A* or *B*) increases with the number of subjects employing this option; and *iv*) an evasion game (EV), where the payoff of either option (*A* or *B*) decreases with the number of subject choosing this option (see Appendix, sections 1 and 2, for details). The best choice context corresponds to a non-social (technological) decision situation, while the other three contexts

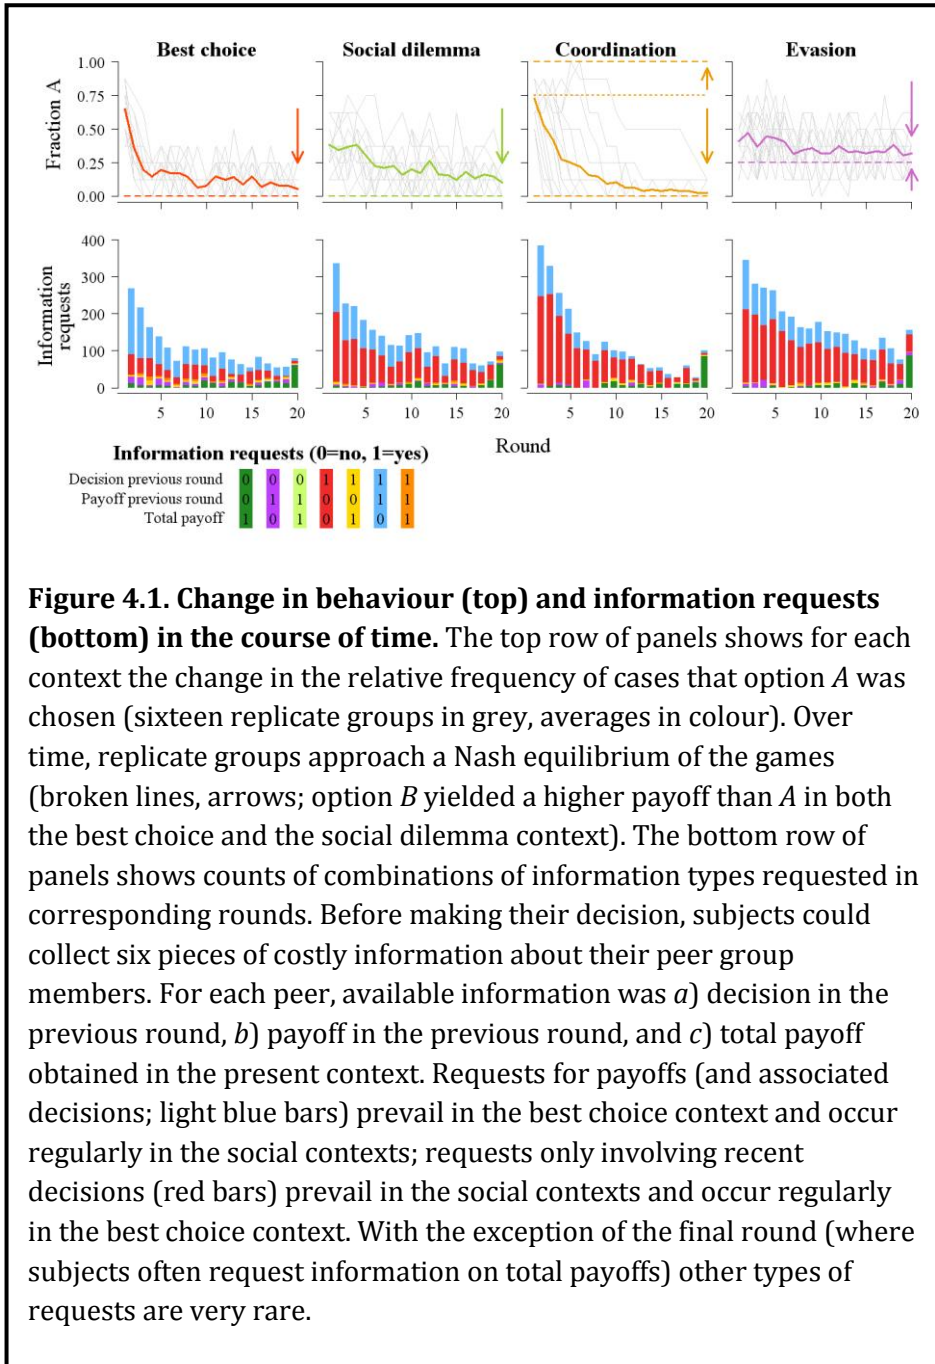
represent different types of social settings. In each context, groups of eight individuals were randomly formed, and subjects interacted for 20 rounds within the same group. All individuals were informed beforehand which type of context they encountered, but they did not know the payoffs associated with the two options. Payoffs were noisy, so that it would require several trials to find out which option was better (see Methods Summary for details). In each round, subjects had to decide simultaneously and anonymously for one of the two options. Before making their decisions, subjects could pay a small cost to collect information about their peers. The information available about each peer included *a*) the previous decision, *b*) previous payoff, and *c*) the total payoff acquired in the current context. If a subject chose to collect information, she could freely request up to six pieces of information. For example, a subject might request all three types of information of two (out of seven) of her peers, or she might request only the previous decision of six of her peers. At the end of each round, subjects were informed about the payoff of their decision in that round. As a benchmark, we ran four control sessions with sixteen subjects each ($n=64$) where no social information could be collected (see Appendix, section 2).

To check for the replicability and robustness of the results obtained, we also conducted a second experiment (henceforth called the ‘follow-up experiment’), consisting of ten sessions with 20 subjects each ($n=200$). This experiment shared the same general set-up with the primary experiment but differed in various aspects (group size was five instead of eight, individuals were allowed to view only four pieces of information per round, the payoff matrices were slightly different, the stochastic component in the payoffs was increased; see Part B of the Appendix for details). We will first report on the outcome of the primary experiment and subsequently address questions related to robustness and replicability by referring to the results of the follow-up experiment.

Results

Dynamics of behaviour and social information use

In the course of time, the behaviour in all groups converged to a Nash equilibrium of the corresponding context (Figure 4.1, top row, broken lines), indicating that subjects adjusted their behaviour according to the underlying payoff structure (see Appendix, section S4.2, for details). The availability of social information was associated with an increased adoption of the superior option in the best choice context, as well as the adoption of the dominant strategy in the social dilemma game (generalised mixed models with subject nested in group as random factors: $P < 0.001$ and $P = 0.047$ for the BC and SD, respectively; see Appendix, section 3 for statistical analysis). Subjects frequently made use of costly social information (Figure 4.1; bottom row; the fraction of decisions preceded by costly request for information was BC: 0.23, SD: 0.25, CO: 0.21 and EV: 0.31). Information requests decreased over time; presumably because there is less need for subjects to adjust their behaviour once a group has approached equilibrium. Furthermore, in 85% of the cases subjects requested two specific types of information: previous peer decisions only (red bars) or the combination of previous peer decisions and associated payoffs (light blue bars). These types of information requests are consistent with frequency-based learning strategies (such as conformism) and payoff-based learning strategies, respectively. On average, subjects requested social information more often when their own previous decision yielded low payoffs. Also, individuals tended to switch when social information indicated that alternative behaviour yielded higher payoffs (consistent with payoff-based learning). Reactions to information on the frequencies of peers' behaviour depended on the interaction context (see Appendix, section S4.4 for a detailed analysis of when individuals requested social information, and how this information affected decision making). In the best choice context, the proportion of information requests including requests for payoffs (0.61) was significantly higher than in the social contexts (SD: 0.40, CO: 0.26, EV: 0.33; Tukey contrasts between effect estimates of the factor 'context' in a generalised linear model: $P < 0.001$). Subjects requested total payoffs almost exclusively in the last round (dark green, light green, yellow and orange bars); presumably out of curiosity to compare their own performance to that of others.



Individual variation in social learning strategies

Subjects strongly varied in their reliance on social learning. In each context, about 20% of all subjects never requested information (Figure 4.2a). In contrast, some other individuals based most of their decisions on information about their peers. The subjects also differed strongly in the degree to which they requested information on payoffs (Figure 4.2b). While the overall fraction of information requests that included requests for payoffs was intermediate in each context (see above), the underlying individual behaviour was surprisingly extreme. For each context separately, we found that most subjects can be categorized in two largely distinct groups with consistent behaviour: those who rarely include payoff information in their requests, and those who do this in the majority of cases (Figure 4.2b).

Consistency of social learning strategies

For each of the four contexts, we categorised subjects' reliance on social learning as low or high, when they requested social information in less or more than 25% of the cases, respectively. This cut-off point divided the data in two roughly equal portions for each context (*cf.* Figure 4.2a). Individuals were significantly more consistent in their reliance on social information (*i.e.*, either categorised as 'high' or as 'low' in all contexts) than expected based on independence between contexts (Figure 4.3a).

Subjects showed a similar consistency with regard to the type of social information they relied on; many of them either strongly relied on payoff-information across all contexts (Figure 4.3b, red bar), or hardly relied on payoff-information at all (Figure 4.3b, blue bar; an individual's reliance on payoffs was categorised as 'low' or 'high', depending on whether her information requests included payoffs in less or more than 50% of the cases). The abundance of these consistent individuals was much higher than expected based on independence between contexts (subjects with consistent high and low reliance on payoffs occurring 6.13 and 2.77 times more than expected on the basis of independence, respectively; $\chi^2=9.811$, $d.f.=1$, $P=0.002$, and $\chi^2=7.093$, $d.f.=1$, $P=0.008$).

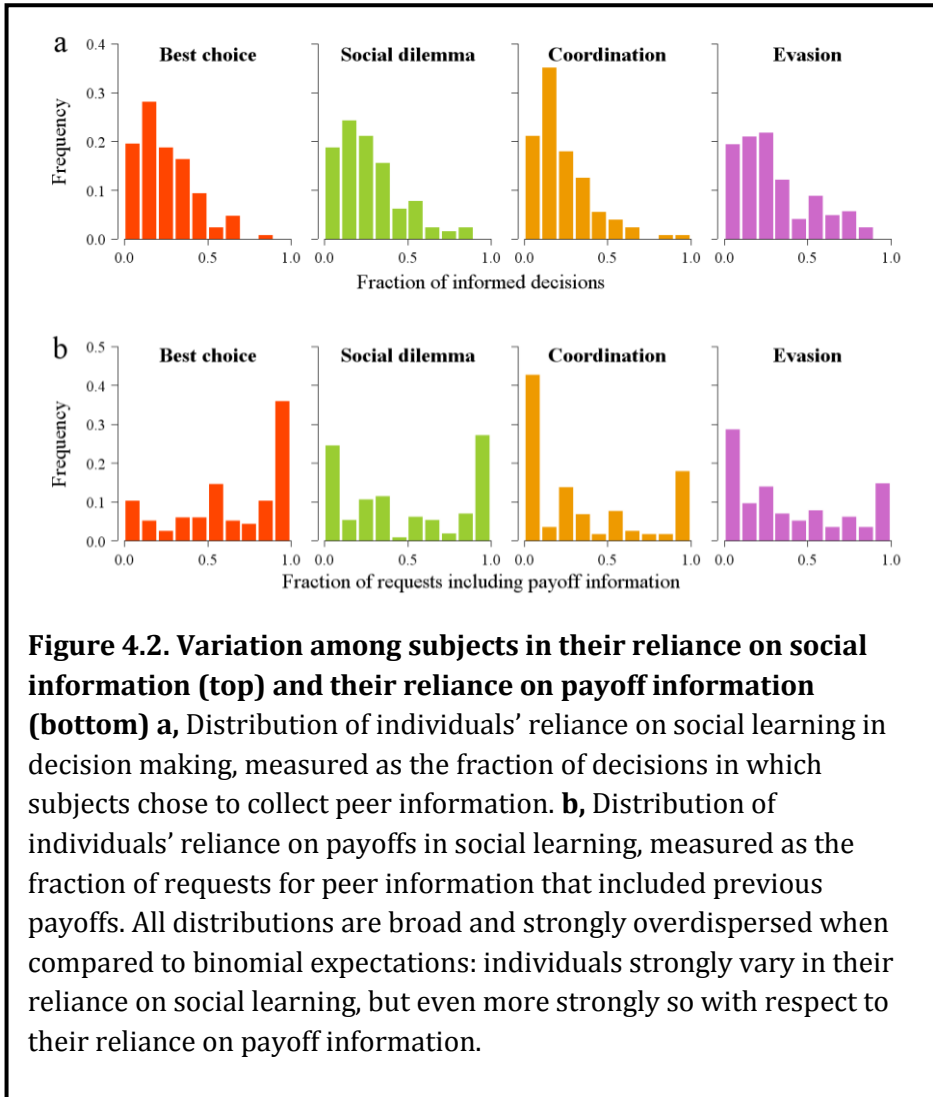
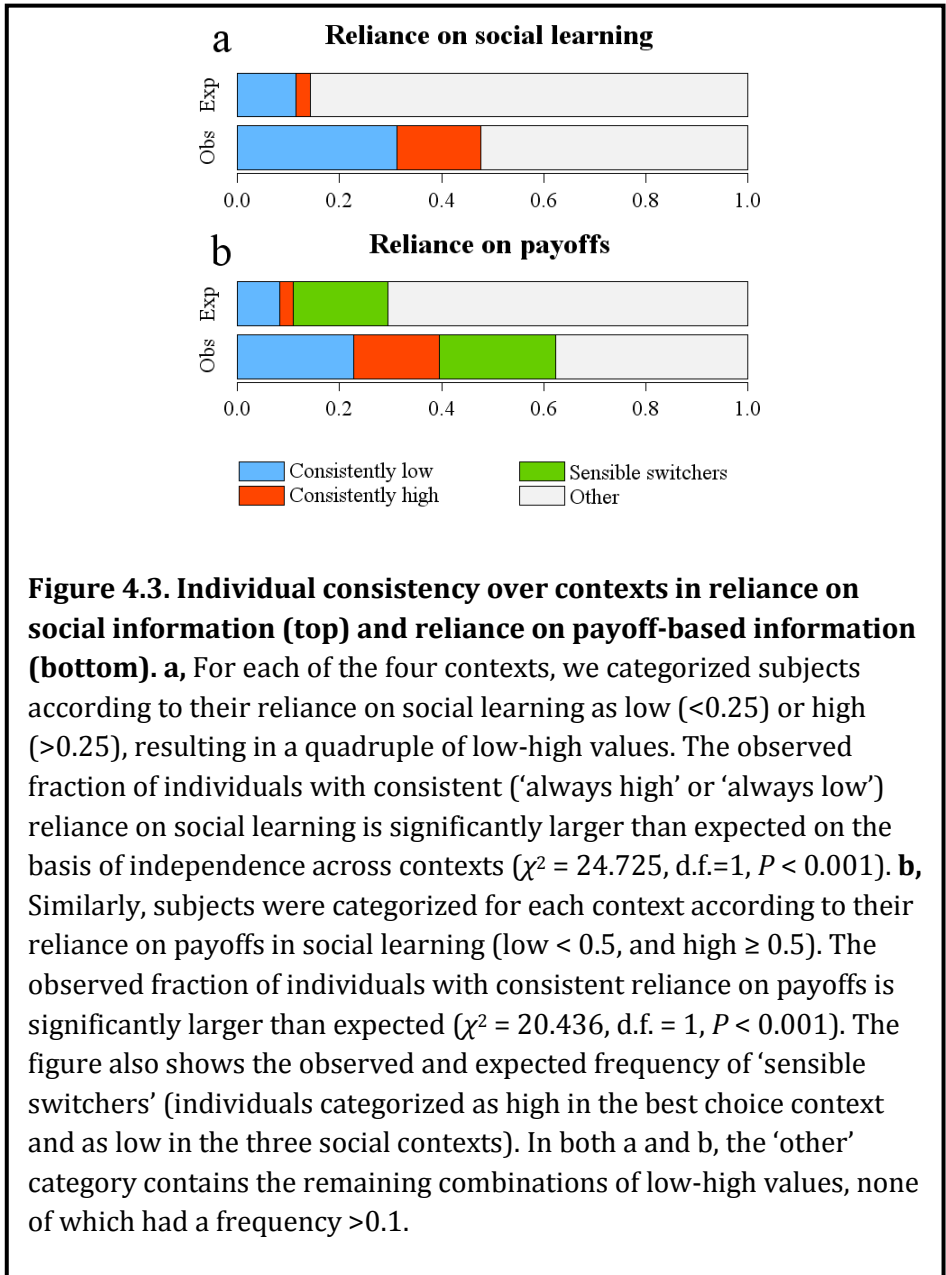


Figure 4.2. Variation among subjects in their reliance on social information (top) and their reliance on payoff information (bottom) a, Distribution of individuals' reliance on social learning in decision making, measured as the fraction of decisions in which subjects chose to collect peer information. **b,** Distribution of individuals' reliance on payoffs in social learning, measured as the fraction of requests for peer information that included previous payoffs. All distributions are broad and strongly overdispersed when compared to binomial expectations: individuals strongly vary in their reliance on social learning, but even more strongly so with respect to their reliance on payoff information.

Another large fraction of subjects employed a specific flexible social learning strategy, with high reliance on payoffs in the (non-social) best choice context and low reliance on payoffs in the other (social) contexts – this is the social learning behaviour that one might intuitively expect (see introduction). However, these 'sensible switchers' are not more abundant than expected based on independence between contexts ($\chi^2=0.319$, d.f.=1, $P=0.572$). We did not observe significant differences between types of social learning strategies and performance in the experiment, as reflected

in total earnings (Tukey contrasts on a linear model: $P > 0.121$ for comparisons between the three types of social learners highlighted in Figure 4.3b).



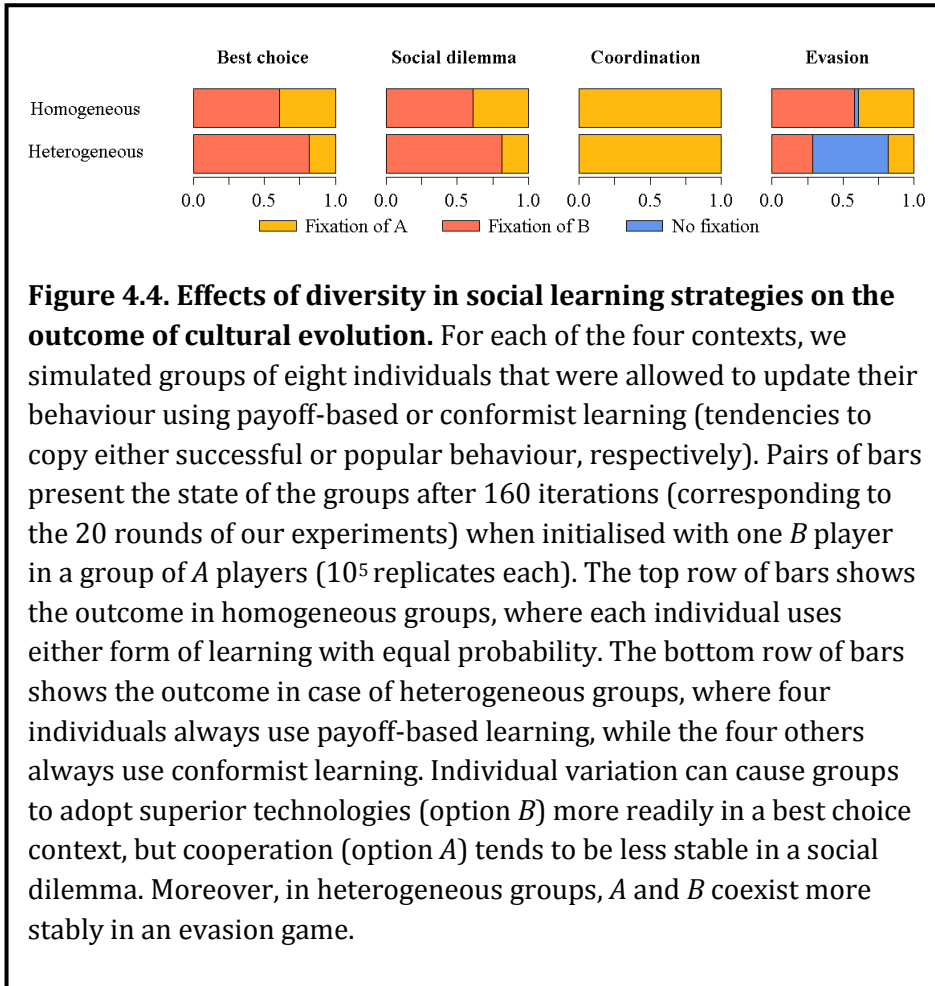
Robustness and replicability

Our follow-up experiment confirmed all main results of the primary experiment (Part B of the Appendix; compare Figures S4.7.1, S4.7.2 and S4.7.3 with Figures 4.1, 4.2 and 4.3 of the main text). On average, the increased magnitude of the stochastic component on payoffs increased subjects' reliance on social information in each of the interaction contexts (ANOVAs; BC: $P=0.035$; SD: $P<0.001$; CO: $P=0.071$; EV: $P<0.001$), which is in line with results reported elsewhere (*e.g.*, Toelch et al. 2009; Morgan et al. 2012). As in the primary experiment, subjects strongly varied in the extent to which they based their decisions on social information. In each of the four interaction contexts, individuals' reliance on payoff information strongly varied; some always included payoff information in their requests, whereas others ignored payoff information altogether, and these differences were even more pronounced than in our primary experiment. Moreover, individuals that were consistent in their social learning strategies across contexts were observed in much higher frequencies than expected based on independence between contexts. In other words, as in our primary experiment, many individuals were consistent in either strongly or weakly relying on social information, and many individuals were consistent in the type of social information they requested.

Implications of individual variation for cultural evolution

Our experimental results indicate that there are strong and consistent individual differences in social learning strategies. But does this individual variation matter? To address this question, we developed a simple model of cultural evolution in which individuals interact in small groups and are allowed to update their behaviour by either payoff-based learning or frequency-based learning. We considered two implementations, a stochastic model and an individual-based simulation. Here, we report on the results of the individual-based simulation; for the results of the stochastic model, see Appendix (section S4.9). We imposed two conditions: a homogeneous population, in which all individuals updated their strategy by payoff-based learning or frequency-based learning with a 50-50 probability, and a heterogeneous population, in which half of the individuals in a group always updated based on payoffs, and the other half always updated based on frequencies. In our simulations, we documented

the fixation probability of behaviour A or B after the introduction of a single B mutant in a population of A players. Figure 4.4 shows the outcome of simulations after an average of 20 updates per individual (in correspondence to the number of rounds in the experiments). In three out of four contexts, the existence of individual variation in social learning strategies considerably affects the outcome of cultural evolution. In the coordination game, payoff-based and frequency-based learning both disfavour the spread of rare behaviours, so B will be rapidly lost both in homogeneous and heterogeneous populations. In the best choice context, a newly introduced superior technology (B) is more likely to reach fixation in heterogeneous populations; payoff-based learners readily adopt B , and the increase in frequency of B reduces the risk that the superior technology disappears due to frequency-based learning or stochastic processes (*e.g.*, copying errors). The same logic applies to the social dilemma: a newly introduced ‘cheater’ mutant will obtain higher payoffs, and payoff-based learners in a heterogeneous population will readily switch to defection; once cooperators are no longer in the majority, conformists start defecting, and cooperation is lost. In the evasion game, variation in social learning strategies enhances the stability of the coexistence of A and B . Interestingly, in heterogeneous populations, social learning strategies become correlated with behaviour in this game (playing A or B ; see Appendix, section S4.8, for model details and additional model analyses). The emerging role differentiation can thus increase average payoffs in a group.



Discussion

Cross-cultural experiments have established that human social behaviour strongly varies between cultures (Henrich et al. 2001; Herrmann et al. 2008). This suggests that the behaviour of typical ('WEIRD'; Henrich et al. 2010) participants in decision making experiments is likely to underestimate the variation present in the human population. Our experiment shows that even within such a relatively homogeneous sample of subjects there is strong and consistent variation in the way people behave. The pronounced individual differences in social learning strategies we report emphasize that the average behaviour in a group or

population is far from representative for the behaviour of its constituent individuals. Furthermore, our simulation model suggests that individual variation can affect the dynamics and outcome of social interactions. Although individual differences have been at the heart of personality research in psychology, they are often disregarded in many other areas of the social sciences. Our results strongly suggest that individual differences need to be taken seriously, both in theoretical models and in the analysis of empirical data. More specifically, we advocate that models of cultural evolution take individual differences in social learning strategies into account.

In the best choice context, the availability of social information was associated with an increase in the adoption of the optimal behaviour. This is in line with earlier findings in the experimental and theoretical literature that social learning can increase an individual's performance in non-social contexts (Boyd and Richerson 1985; Efferson et al. 2008; Rendell et al. 2010). In agreement with earlier experimental evidence (McElreath et al. 2005; Efferson et al. 2008), also in a non-social context subjects differed substantially in the way they used social information. These differences were even more pronounced in the social contexts. Recent experimental results suggest that variation in social learning strategies might be associated with variation in behaviour in a given context (manuscript in preparation). For example, it is plausible that payoff-based learning in a social dilemma will be associated with a higher tendency to defect (since this is the payoff-dominant strategy), while frequency-based learning (conformism) might stabilize cooperative behaviour (Boyd and Richerson 1985; Molleman et al. 2013a; Molleman et al. 2013b).

Of course, the results of highly stylized experimental study like ours should not be over-interpreted. Decision making experiments in a lab setting allow for a high degree of control, but inevitably abstract from potentially important aspects of reality. Even though our experimental design imposed relatively few restrictions on how individuals could track their social environment, our study does not account for a variety of factors that play a role in real-life cultural transmission. For instance, it was impossible for subjects to bias their attention towards older, more experienced, or more dominant peers (Dugatkin and Godin 1993; Duffy et

al. 2009; Henrich and Broesch 2011), and active teaching (Hoppitt et al. 2008) was excluded by design. Also, subjects in our experiment were informed about the general structure of the interaction contexts in which they were interacting. In reality, this structure is typically unclear and people have to base their decisions on their best guess of how their behaviours and the behaviour of others affect their payoffs. More specific to our design including different contexts, subjects in a session always encountered the simplest (non-social) best choice situation first and subsequently interacted in the other (social) contexts in random order. Some of these concerns can be addressed by additional experiments. For example, a recent experiment revealed that the above-mentioned order of contexts encountered did not affect the results; presenting the coordination game first yielded almost identical results concerning social information use (in preparation). Still, lab experiments should best be viewed as unravelling interesting features of human behaviour that should subsequently be scrutinized under more realistic conditions.

Assuming that the individual variation found in our study is 'real', why should individuals differ in their learning strategies, and why should they be consistent in their use of social information across different contexts? Possible explanations can be found in the literature on 'animal personalities'. In recent years, evidence has accumulated that consistent individual differences in behaviour ('behavioural syndromes' or 'personalities') are not only prevalent in humans, but also exist throughout the animal kingdom (Gosling 2001; Sih et al. 2004). In the past, it has often been assumed that selection on behaviour should result in a single, optimally adapted phenotype. By studying animal personalities it has become apparent that a variety of mechanisms (reviewed in Wolf and Weissing 2010) can lead to the evolutionary emergence of consistent individual variation in behaviour. More specific models based on this theoretical framework predict the coexistence of alternative life history strategies (Wolf et al. 2007; Réale et al. 2010); the coexistence of individuals differing strongly in responsiveness and social sensitivity (Wolf et al. 2008; Wolf et al. 2011); and the coexistence of communication strategies that differ in the way that signals are sent and how information is processed and interpreted (Botero et al. 2010). These results may serve as a useful starting point for understanding the consistent individual

differences in social learning strategies observed in this study. For example, Botero and colleagues (2010) have shown that a ‘general-purpose’ mechanism is often not able to produce optimal behaviour in specific situations; this leaves opportunities that can be exploited by other strategies, allowing for coexistence. These insights were obtained in the context of communication, but they might apply to learning as well: learning strategies are also applied in a multitude of contexts, and are unlikely to be optimally tailored to every particular learning situation. Cognitive constraints and incomplete information on the learning context will often prevent the evolution of learning strategies that are optimal in every possible circumstance. Instead, evolution may result in the coexistence of locally suboptimal but cognitively ‘cheap’ strategies using rules-of-thumb (*cf.* Gigerenzer et al. 1999) such as payoff-based learning or frequency-based learning. Other mechanisms facilitating the coexistence of behavioural strategies may apply to social learning as well. For example, the benefits of information are often negatively frequency-dependent: the value of information decreases when more individuals also have this information. This immediately explains why individuals relying heavily on information coexist with others that are seemingly not interested in gathering information at all (Wolf et al. 2008). Last, but not least, the coexistence of learning strategies could be explained by ‘synergy’ between different forms of learning. As suggested by our simulation model, group heterogeneity with respect to learning strategies may under some circumstances lead to more efficient outcomes. Such synergetic effects might again induce negative frequency-dependent selection and result in the coexistence of alternative learning strategies.

Even if different learning strategies do coexist, one might expect that individuals are flexible and employ different strategies in different contexts. Again, the literature on animal personalities has revealed various mechanisms explaining a more rigid behaviour, corresponding to the consistent use of the same type of behaviour in different contexts. Behavioural consistency can be favoured when strategic conventions are establishing (McNamara and Weissing 2010) or when behaviour serves as a signal to others. For example, already a small fraction of ‘socially responsive’ individuals in a population may exert a strong selection pressure in favour of consistency (Wolf et al. 2011). Consistency can also

be supported by all kinds of positive feedbacks: individuals can increase their performance with the consistent use of a certain social learning strategy, due to an increased efficiency in the collection, interpretation and application of the kind of information as they get more experienced in using this strategy.

Rather than explaining individual variation from an evolutionary perspective, scholars of human personality have focused on characterising the structure of individual differences. This has resulted in a number of models that aim to describe human personality along a few dimensions (the 'Five Factor Model' (Digman 1990) being best-known example). It is not unlikely that the observed variation in social learning strategies is associated with these dimensions or factors, and this possibility presents an interesting objective for further research. For instance, one might predict that reliance on social information is associated with openness, or with an orientation towards collectivism (as opposed to individualism (Oyserman et al. 2002)); in fact, this latter association has just recently been observed by Toelch and colleagues (personal communication). Similarly, one might expect that individuals that focus on payoff-information are inclined to behave more competitively (rather than cooperatively) in social interactions.

Another interesting future direction would be to address whether the presence of individual differences in social learning strategies affect the dynamics of behaviour (as predicted by our simulation model). One could experimentally test this by creating groups of individuals assorted according to their social learning strategies, and comparing these homogeneous groups to unassorted (heterogeneous) groups. Another interesting question that could be tested in such a setup is whether the behavioural dynamics are different between groups that are homogeneous for different social learning strategies (such as payoff-based learning and frequency-based learning). For example, one might expect that a group of payoff-based learners may more readily adopt new superior technologies, whereas groups of frequency-based learners may be faster in reaching equilibrium in a coordination context.

Methods

Experimental sessions were run at the Sociology department of the University of Groningen. Subjects (mainly students, aged 19-25) participated in twenty separate sessions (primary experiment: eight sessions; control: four sessions; follow-up experiment: ten sessions), each lasting for about 90 minutes. In each session of the primary experiment, sixteen subjects participated. Subjects were paid according to their performance in the experiment, and earned €29.60 on average, excluding a show-up fee of €7. All subjects interacted in four blocks of 20 rounds in sequence, each representing a different interaction context. At the beginning of each block, subjects were randomly grouped, labelled with a number 1-8, and each received 3,000 points to play with (1,000 points = 1 euro). These points could be used to collect social information. During the block of rounds subjects could make substantial profits, but subjects could also lose their initial endowment in case of negative outcomes. After the groups were formed, participants read brief instructions on their screen and filled out a quiz checking understanding of the upcoming block of rounds. In the primary experiment, sessions started with the best choice context; the other contexts were played in random order.

In each of the contexts, subjects repeatedly decided between two options. The setting was framed as planting either of two crops on a farm. Payoffs for option A or B were $p \cdot a + (1 - p) \cdot b$ and $p \cdot c + (1 - p) \cdot d$, respectively, where p denotes the fraction of A players in the group. We chose the following parameter settings: best choice $\{a=50, b=50, c=300, d=300\}$, social dilemma $\{a=300, b=-250, c=600, d=0\}$, coordination game $\{a=175, b=-75, c=-75, d=675\}$ and the evasion game $\{a=-75; b=175; c=675; d=-75\}$. A stochastic component was added to each of the subjects' payoffs independently, by adding a number taken from a standard normal distribution with mean 0 and standard deviation 175. Sessions ended with a questionnaire including items on game play and game understanding. The experiment was conducted using Z-Tree (Fischbacher 2007); code is available upon request. See Appendix, section 6, for experimental instructions and screenshots.

Simulations were programmed in C++ (code available on request). We tracked the cultural evolution of behaviour after introducing a single B

player in a group of A players. Each time step of the simulation, a random individual ('focal') is drawn from a group of eight to update its behaviour by either payoff-based or frequency-based learning. This process is iterated for 160 time steps, so that each individual has on average been drawn 20 times. Social learning strategies correspond to strategies feasible in our experiment: in frequency-based learning the focal switches if the other behaviour is more abundant in a sample of six others. In payoff-based learning, the focal switches to the behaviour of the individual with the highest payoff in a sample of three others.

Appendix

Part A. Primary experiment: additional information and analyses

In this Part we give an overview of the experimental setup of our primary experiment, and provide additional data analyses. First, we delineate the general setup and describe the four contexts in which the subjects interacted. Second, we present the results of control sessions of our primary experiment, in which no peer information was available to the subjects, so that social learning was impossible. Third, we present additional analyses of the data from the primary experiment. We focus on the question under which conditions individuals used social information, and how this information affected decision making. Fourth, we present regression models fit to the distributions of the social learning strategies presented in Figure 4.2 of the main text.

S4.1. General setup

We confronted subjects with four different interaction contexts in sequence. Decision making was formulated as planting crops on the land of an imaginary farm. Each interaction context consisted of a block of twenty rounds, in which subjects could decide which of two crops they would plant on their farm (we used actual crop names, but for simplicity we will refer to the options as *A* and *B*). At the beginning of each block, subjects were randomly divided in two groups of eight and informed about how the payoffs in the current context depended on their own decisions, and the decisions of their group members. Before a block of rounds started, subjects filled out a brief quiz to check their understanding of the structure of the coming interaction context.

In each round, all subjects made their decisions simultaneously. After making their decision, subjects were each informed about the payoff they obtained in that round. Before making their decisions, subjects could choose whether or not to collect information about the members of their group (they had to pay a small cost for this). If they did so, they could view a maximum of six pieces of social information. For each group member,

the information available included *a*) the previous decision, *b*) the previous payoff and *c*) the total payoff accumulated in the current context. After a block of twenty rounds ended, new groups were formed, and subjects were informed about the payoff structure of the new context (again followed by a short quiz).

Each of the interaction contexts was associated with a payoff matrix $\begin{pmatrix} a & b \\ c & d \end{pmatrix}$. Let p denote the fraction of subjects in a group that chose option

A. The deterministic component of the payoff of a subject that chose *A* was given by $p \cdot a + (1 - p) \cdot b$. Similarly, for a subject that chose *B* this was given by $p \cdot c + (1 - p) \cdot d$. Payoffs were noisy, reflecting the fact that the outcome of behaviour is often influenced by exogenous factors. This noise made it harder for the subjects to find out individually which choice was optimal, and as a consequence, made social information more valuable. Noise was implemented by adding a stochastic term to the deterministic component of the payoffs of each subject separately. This stochastic component was a number drawn from a normal distribution with mean 0 and standard deviation σ .

S4.2. Interaction contexts

In this section, we describe the four interaction contexts, their associated payoff matrices, and the Nash equilibria of the one-shot version of each context, corresponding to the dotted lines of Figure 4.1 of the main text. Subjects were always confronted with the ‘best choice’ context first. Starting with this most simple interaction context facilitated the subjects’ understanding of the context they were in and how their decisions affected their payoffs. The three other interaction contexts – in which payoffs depended on the decisions of others – were played in randomized order. These other contexts correspond to the three different classes of games: a social dilemma, a coordination game and an evasion game. The payoff matrices of each of the interaction contexts were chosen such that at $p=0.5$, the payoff difference between *A* and *B* equalled 250 points (1000 points=1 euro). We set $\sigma=175$, so that at $p=0.5$, choosing the option with

the lower expected payoff nevertheless lead to a higher payoffs in 16% of the cases.

In the *best choice* context (BC), a subjects' payoff did not depend on the decisions of the other subjects in their group ($a=b$ and $c=d$); the payoff matrix was given by $\begin{pmatrix} 50 & 50 \\ 300 & 300 \end{pmatrix}$. Irrespective of the behaviour of the other players, option *B* yielded higher payoffs on average. In game theoretical terms, *B* dominates *A* and therefore $p^*=0$ is the Nash equilibrium of the one-shot version of the 'game' (technically, it is not a game, since the payoffs do not depend on the actions of others).

In the *social dilemma* (SD), cooperation (*A*) is dominated by defection (*B*): $a < c$ and $b < d$; the payoff matrix was given by $\begin{pmatrix} 300 & -250 \\ 600 & 0 \end{pmatrix}$. Therefore, like in the best choice context, $p^*=0$ was the Nash equilibrium of the one-shot version of the game. However, all subjects obtain a higher payoff when they all cooperated ($p=1$), compared to this equilibrium ($a > d$). This shows that in this context, collective interests and individual interests are opposed to each other, like in the famous (two-player) Prisoner's Dilemma game.

In the *coordination game* (CO), the payoff of choosing one of the two options increased with the number of others also choosing it ($a > c$ and $b < d$). The payoff matrix was given by $\begin{pmatrix} 175 & -75 \\ -75 & 675 \end{pmatrix}$. In this case, both $p^*=0$ and $p^*=1$ are Nash equilibria. However, the Nash equilibrium at $p^*=0$ is Pareto superior; the payoff to all players was higher compared to the equilibrium at $p^*=1$. There was a (dynamically unstable) equilibrium located at

$$p^* = \frac{d-b}{a-b-c+d}. \quad (\text{S4.1})$$

In this equilibrium (in our case, when two subjects chose *A* and six subjects chose *B*), the expected payoffs of choosing *A* and choosing *B* were

equal, but any deviation from this equilibrium leads towards either of the Nash equilibria. Typical examples of coordination games are the Stag-hunt game and the Battle of the Sexes.

In the *evasion game* (EV), the payoff of choosing one of the two options decreased with the number of others also choosing it ($a < c$ and $b > d$).

The payoff matrix was given by $\begin{pmatrix} -75 & 175 \\ 675 & -75 \end{pmatrix}$. In this case, there was a

Nash equilibrium given by equation S1 (when two subjects choose *A* and six subjects choose *B*). Typical examples of evasion games are the Hawk-Dove game and the Snowdrift game.

S4.3. Control sessions

To investigate how the presence of social information affected behaviour, we ran control sessions in which subjects could not collect information from their peers before making their decisions. Otherwise, the setup was identical to the setup of our primary experiment. We ran four independent sessions with sixteen subjects each ($n=64$), for a total of eight replicate groups of eight participants each. Figure S4.3.1 summarises the development of the decisions in the twenty-round blocks of the four interaction contexts, and allows for comparison to the results presented in the main text. Comparing the dynamics over time to the dynamics in our primary experiment (where social information was available; cf. Figure 4.1 of the main text), we observe differences in the best choice context and the social dilemma. In both of these contexts, groups tended to approach the Nash equilibria of the one-shot versions of the game (in both games characterised by $p^*=0$) more slowly in the control sessions, compared to the primary experiment (two separate binomial generalized linear mixed models (GLMM) with subject nested in group as random factor, and with ‘information present/absent x round’, ‘previous decision x previous payoff’ and ‘difference between payoffs in round $t-1$ and $t-2$ ’ as fixed factors, detected significant effects of the presence of information on the probability that an individual chooses option *B*; $P<0.001$ and $P=0.047$ for BC and SD, respectively). This suggests that the availability of social

information allowed for a faster adoption of individually optimal behaviours. The presence of social information led to a more rapid convergence on the superior option in the best choice context, but it reduced cooperative behaviour in the social dilemma. In the coordination and evasion game, fitting such a statistical model to data is not as straightforward, since the ranking of payoffs ($A < B$, or $A > B$) of decisions depends on the group frequency of A and B .




Figure S4.3.1. Decisions in the absence (top) and presence (bottom) of social information. The top row of panels shows frequencies of choices for option A over twenty rounds of the four interaction in four control sessions, where social information was not available (eight replicate groups in grey, averages in colour). As a reference, the bottom row of panels shows the decision data from our primary experiment (where social information was available), also presented in Figure 1 of the main text. Over time, groups tended to approach the Nash equilibria of the one-shot versions of the games (indicated by the arrows and the dashed lines).

S4.4. The effects of social information on decision making

In the main text, our analysis of social learning strategies focuses on the frequency and the types of social information that individuals requested before making a decision. However, these factors constitute only one aspect of an individual's social learning strategy; a more complete account would include how this information affected decision making. This section aims to look deeper into how social information affected decision making in our primary experiment. First, we analyse how subjects' decisions to collect social information depended on the payoffs they received in the previous round (Figure S14.1). In line with earlier findings (Morgan et al. 2012), we find a strong negative relationship between the previous payoff and the request rates for social information. Second, we consider how social information affected decision making in each of the interaction contexts. To this end, we separately consider decisions that were preceded by requests for others' previous decisions *and* previous payoffs (cf. payoff-based learning; Figure S4.4.2), and decisions that were preceded by requests for (only) others' previous decisions (cf. frequency-based learning; Figure S4.4.3). We find strong indications for payoff-based learning in each of the four interaction contexts. We also find indications of frequency-based learning in each of the contexts, but the effects of observing only the decisions of group members are less pronounced.




Figure S4.4.1. Information request rates as a function of individual payoff in the previous round. Data on previous payoffs were pooled in cohorts of 100 points to obtain the symbols reflecting request rates. Error bars represent 1 standard error of the mean (SEM). Total number of observations per context was 2432. Cohorts with fewer than five observations are not shown as a data point. We fitted a binomial GLMM to decisions to request information, using ‘subject’ as random effect, and ‘round’ and ‘interaction context’ as fixed factors. This analysis detects a negative effect of previous payoffs ($P < 0.001$) on the probability to request information, indicating that subjects tended to request social information when their current behaviour had unfavourable returns. This GLMM also detects that rates of requests for social information decrease over time ($P < 0.001$), confirming the observations of Figure 4.1 of the main text.


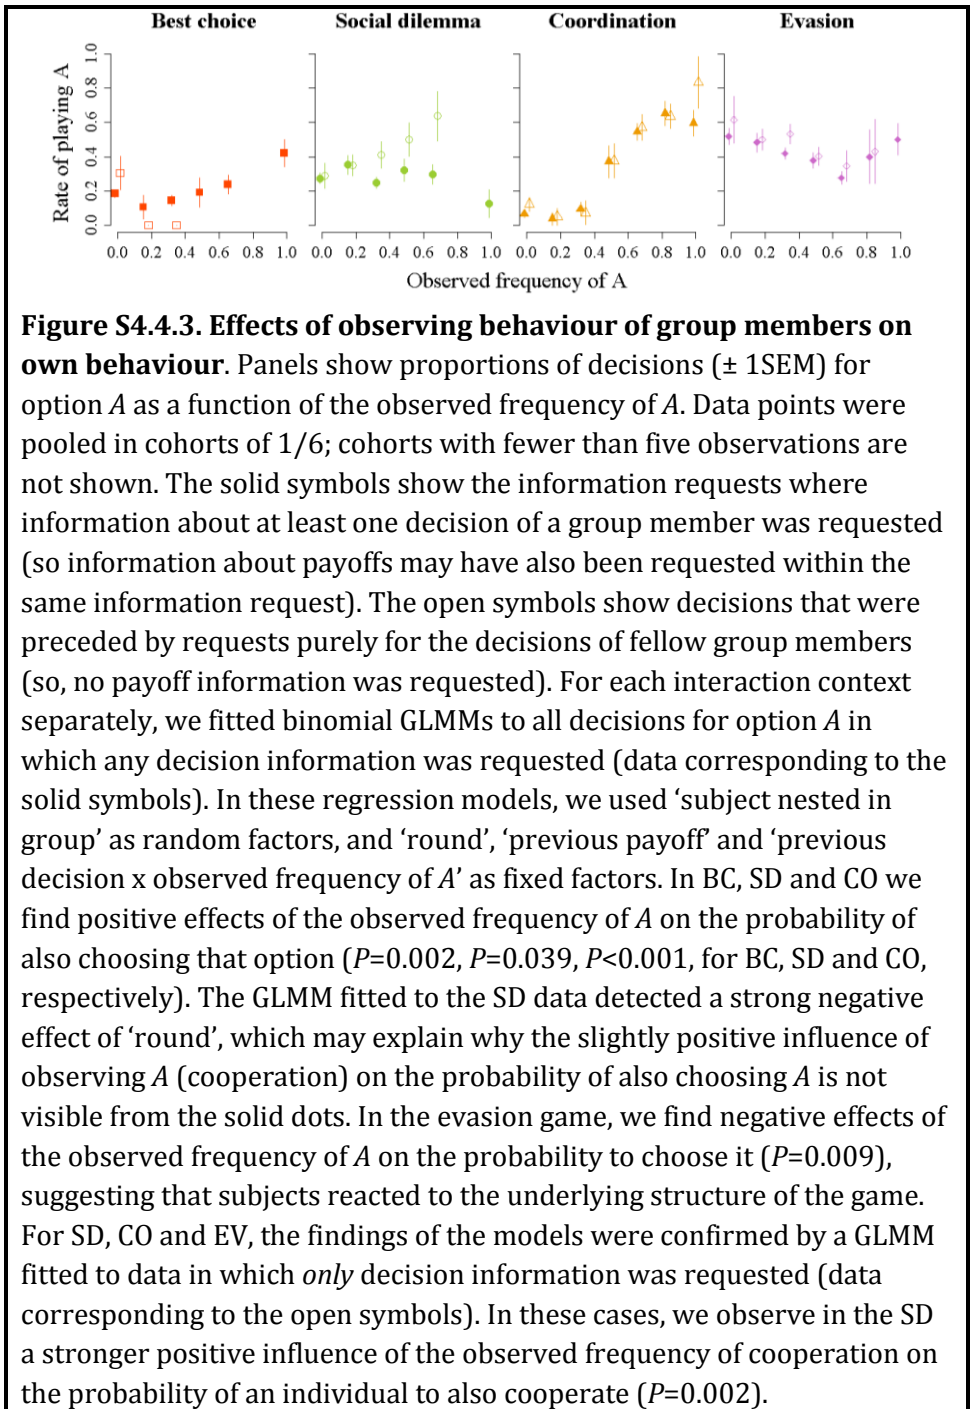


Figure S4.4.2. Effects of payoff information on switching

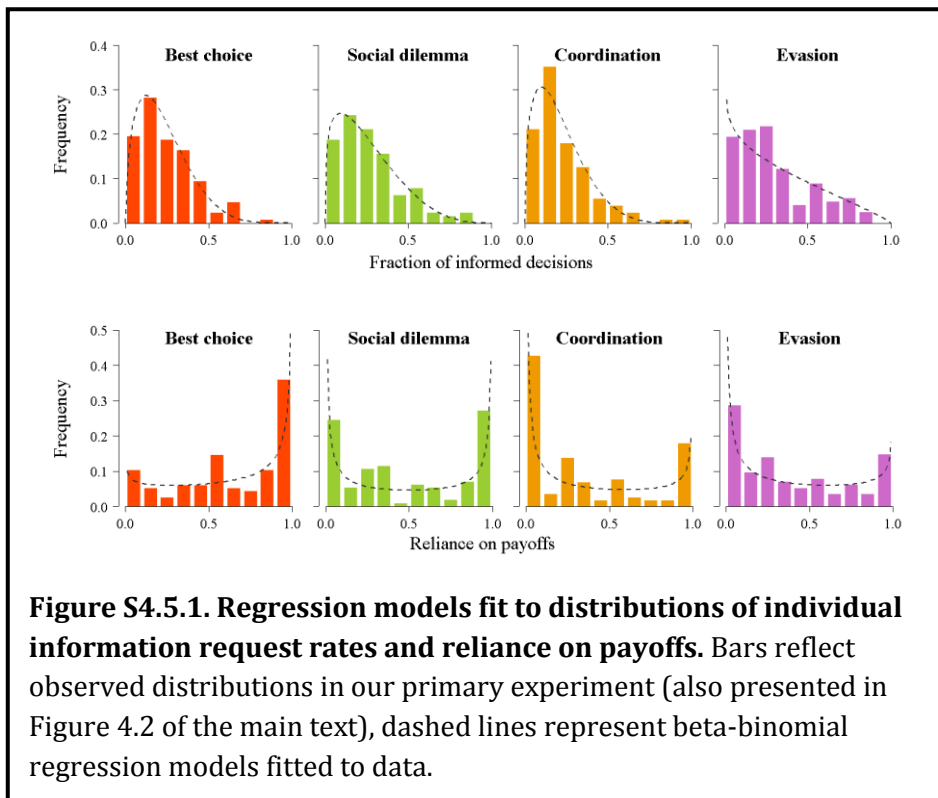
behaviour. Panels show proportions of decisions to switch behaviour as a function of the difference in own previous payoff and the observed average payoff of the alternative option. Open and solid symbols represent rates of switching ($\pm 1\text{SEM}$) from *B* to *A* and from *A* to *B*, respectively. Payoff differences were pooled in cohorts of 200 points to obtain the data points shown (cohorts with fewer than five observations were omitted). We fitted a binomial GLMM to decisions to switch to each context separately, with subject nested in group as random factors, and ‘round’ and ‘previous decision x previous payoff x observed payoff difference’ as fixed factors. Each of these models detected significant effects for the observed payoff difference on the probability of switching ($P < 0.001$ in BC, CO and EV, and $P = 0.023$ in SD). Also, an overarching GLMM with the same error structure fitted to decisions to switch for of all contexts, in which ‘interaction context’ were added to the interaction terms, detected an overall significant effect ($P < 0.001$) for the observed payoff difference. These findings are strong indications for payoff-based learning.



S4.5. Regression fit to social learning strategy distributions

Figure 4.2 of the main text shows the distribution of social learning strategies observed in our experiment. Here, we further analyse those distributions by fitting regression models to this data in order to characterize patterns of variation. For each of the interaction contexts separately, we fitted beta-binomial models to *i*) individual rates of information requests (cf. reliance on social learning) and *ii*) proportion of information requests which included previous payoffs (cf. reliance on payoffs). These measures are analysed as the probability of success in Bernoulli trials, where a ‘success’ refers to a decision to request information, and a request including payoff information, respectively. The beta-binomial distribution has two positive parameters (α, β) , which can produce a wide range of distributional forms. For instance, $\alpha=\beta=1$ results in a discrete uniform distribution (where each number of successes out of n trials is equally likely), and for large α and β , the binomial distribution is approached. When $\alpha < 1$ and $\beta < 1$, the distribution is U-shaped.

Model selection was conducted by maximum likelihood estimation of our observed data for alternative beta-binomial distributions (defined by parameters α and β) using the ‘optim’ function in *R*. The models with the lowest log-likelihood values are shown in the dotted lines of Figure S4.5.1. For each game, the parameters (α, β) of the best-fitting distributions are shown in Table S4.5.1. Rates of informed decisions show a broad, hump-shaped distribution, highlighting individual diversity in reliance on social learning. Individuals differ even more strongly in their reliance on payoffs, with an even broader, U-shaped distribution with peaks at the extremes. Fitting mixtures of beta-binomial distributions yielded better fits (based on their AICs) to the distributions in reliance on payoffs in three out of four interaction contexts. This more complicated analysis lead to the same conclusions. For clarity of presentation, we chose to show the results of the current simpler analysis; details of the more detailed analysis are available on request.



	Reliance on social information		Reliance on payoffs	
	α	β	α	β
Best choice	1.631	5.431	0.786	0.429
Social dilemma	1.283	3.745	0.317	0.320
Coordination	1.492	5.505	0.279	0.510
Evasion	0.864	1.844	0.408	0.620

Table S4.5.1. Parameters of the selected beta-binomial distributions fit to data. Table entries correspond to the dashed lines shown in Figure S4.5.1.

Part B. Follow-up experiment

This Part presents the setup and results from the follow-up experiment. We highlight how decisions and requests for social information developed over time, and how social information use varied between individuals.

S4.6. Setup

To test the robustness of the findings of our primary experiment, we conducted a follow-up experiment consisting of ten independent sessions, each with 20 subjects ($n=200$). The basic setup of this follow-up experiment was the same as the primary experiment reported in the main text: we confronted participants with four interaction contexts in sequence, each consisting of a block of 20 rounds. In each round subjects made a decision between two options, and before making their decision, subjects could collect social information about their peers. However, there were a few differences with the primary experiment: *i*) subjects interacted in groups of five instead of eight; *ii*) because the group size was not the same, the payoff matrices were changed (see below); *iii*) the stochastic component of the payoffs was relatively larger (see below) *iv*) the maximum number of pieces of information subjects could collect in a round was limited to four, instead of six; *v*) subjects could no longer request information about the total payoffs of their group members (this type of information was requested at a low frequency in our primary experiment), but only about their previous decisions and previous payoffs.

The interaction contexts of the follow-up experiment were characterised

by the following payoff matrices: best choice: $\begin{pmatrix} 100 & 100 \\ 300 & 300 \end{pmatrix}$, social

dilemma: $\begin{pmatrix} 400 & -200 \\ 600 & 0 \end{pmatrix}$, coordination game $\begin{pmatrix} 500 & -300 \\ -300 & 900 \end{pmatrix}$, evasion

game $\begin{pmatrix} -300 & 500 \\ 900 & -300 \end{pmatrix}$. As in the primary experiment, we chose the

parameters such, that at $p=0.5$, the payoff differences between A and B are

equal (*i.e.*, 200) in each of the interaction contexts. In the coordination game, the unstable equilibrium was situated where three players choose *A* and two players choose *B*. The Nash equilibrium of the one-shot version of the evasion game was situated where two players chose *A* and three chose *B*. All other equilibria were the same as in the primary experiment. The stochastic component on the subjects' payoffs was relatively larger than in our primary experiment ($\sigma=200$). This implies that at $p=0.5$, choosing the inferior option nevertheless lead to higher payoffs in 24% of the cases.

S4.7. Main results

To allow for comparison with the primary experiment, we present the results along the same lines as in the main text (cf. Figure 4.1, 4.2 and 4.3). Figure S4.7.1 presents the average behaviour in the 20 replicate groups over time (top row of panels), and gives an overview of the types of information requests (bottom row). As in the primary experiment, we find that groups approach the Nash equilibria of the one-shot versions of the games over time in each of the contexts (on which we based the coloured dashed lines and arrows; see section 2 of this Appendix). Comparing these results with Figure 4.1 of the main text, we observe that in the best choice context, dynamics were similar to the primary experiment. In the social dilemma, levels of cooperation tended to be higher than in the primary experiment, possibly due to the smaller group size. In the coordination game, a higher proportion of groups ended up at the equilibrium where all group members choose *A*, which is Pareto inferior to the equilibrium where all group members play *B* (*i.e.*, all players would obtain higher payoffs in the latter equilibrium). This may be explained by the fact that the Pareto inferior equilibrium had a relatively larger basin of attraction in the follow-up experiment than in the primary experiment; the unstable internal equilibrium is closer to 0.5. In the evasion game, groups were attracted to the internal equilibrium (as occurred in the primary experiment).

Information requests tended to decrease over time in each of the contexts (Figure S4.7.1, bottom row of panels). We observe that fewer decisions were based on social information in the best choice context, compared to

the other interaction contexts. Also, we find that most requests were for either a combination of previous decision *and* associated payoffs (blue bars), or for previous decisions only (red bars). In the best choice context, requests for both decision and payoffs are the dominant type of information request. In the other contexts, most requests were for previous decisions only. These observations are in line with those from the primary experiment.

As in the primary experiment, subjects in the follow-up experiment strongly varied in both the amount and the type of information they requested. Figure S4.7.2 shows the distributions of information request rates (reliance on social information, Figure S4.7.2a) and the fraction of information requests that included a request for payoff information (reliance on payoffs, Figure S4.7.2b). The strong overrepresentation of extreme strategies – some individuals always requested payoff information, and others disregard it altogether – confirms the findings presented in Figure 4.2 of the main text.

To assess the individual consistency in social information use across contexts, we analysed the requests for information in the same way as in the primary experiment. For each of the four contexts, we categorized subjects' reliance on social learning as 'low' or 'high' (requesting information in <25%, or >25% of the rounds, respectively). Similarly, subjects' reliance on payoffs in social learning was categorized as 'low' or 'high' (with <50%, or >50% of the requests including payoff information, respectively) for each of the contexts separately. Next, we calculated the expected percentage of consistent individuals, assuming independence of behaviour in the different contexts. Figure S4.7.3 compares those expected fractions of individuals with consistent (high or low) reliance on social information with the observed values (top), and makes the same comparison for reliance on payoff information (bottom). As in our primary experiment, we find that the fraction of individuals that was consistent in their social information use is significantly larger than expected ($\chi^2 = 29.551$, d.f.=1, $P < 0.001$). This is also the case for individuals' consistency in reliance on payoffs ($\chi^2 = 13.673$, d.f. = 1, $P < 0.001$). 'Sensible switchers' (that tended to rely on payoff information in the best choice context only) occur at a high frequency, but not more than expected ($\chi^2 = 2.801$, d.f. = 1, $P < 0.094$).

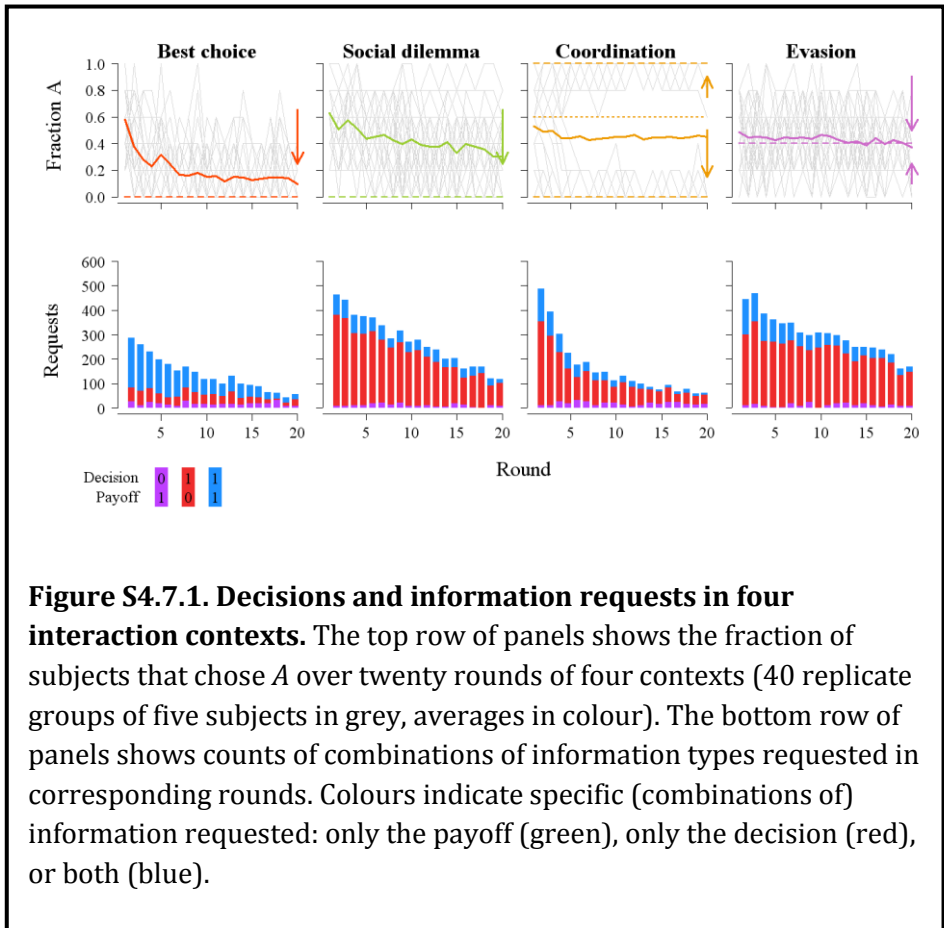
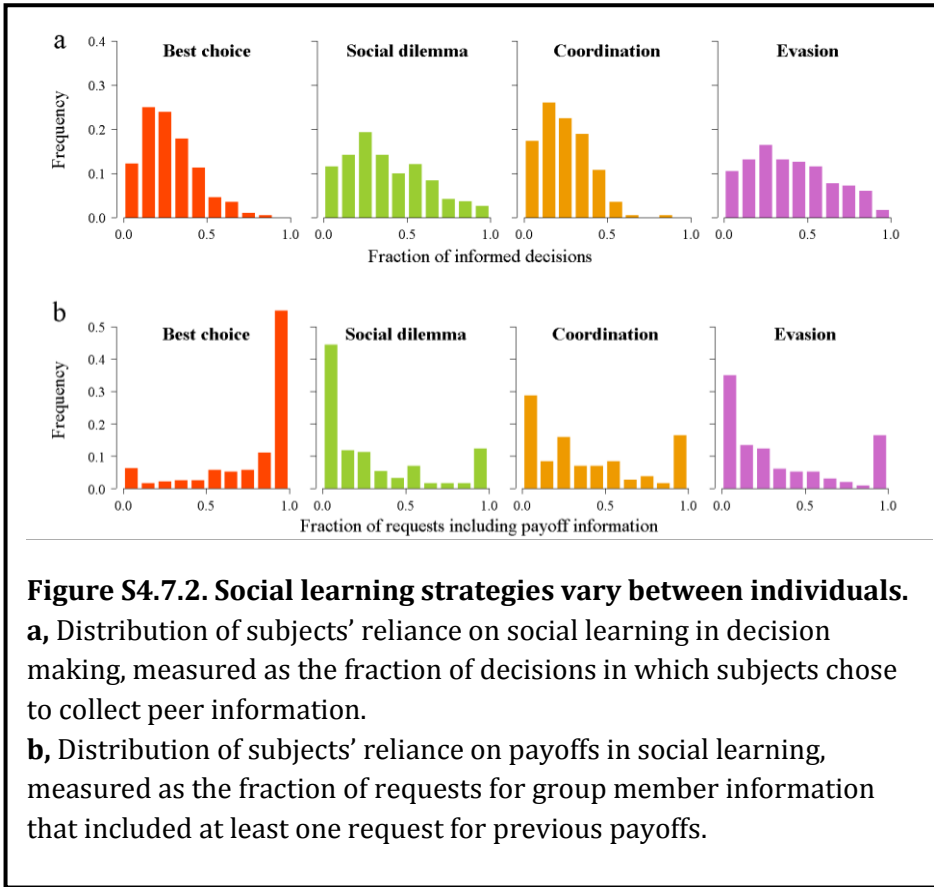
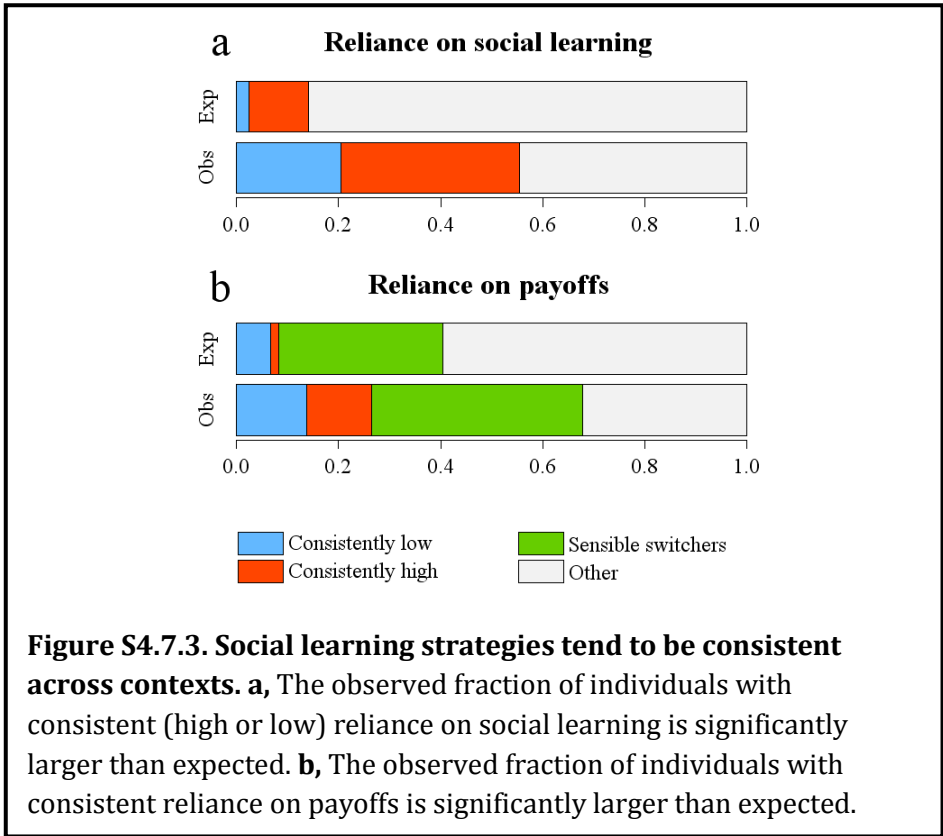


Figure S4.7.1. Decisions and information requests in four interaction contexts. The top row of panels shows the fraction of subjects that chose A over twenty rounds of four contexts (40 replicate groups of five subjects in grey, averages in colour). The bottom row of panels shows counts of combinations of information types requested in corresponding rounds. Colours indicate specific (combinations of) information requested: only the payoff (green), only the decision (red), or both (blue).





Part C. Implications of individual variation in social learning strategies for cultural evolution

In this Part we offer more detailed information about the simulation model presented in the main text. To check the robustness of our findings, we also present an alternative analytical model focusing on cultural evolution of social behaviour in an evasion game. Both models are conceptual ‘toy models’, providing a first theoretical investigation on whether individual variation in social learning strategies influences the dynamics and outcome of cultural evolution. With these models, we do not intend to make any testable predictions or accurately mimic reality. Rather, we aim to provide a proof of principle that individual differences in social learning strategies potentially have a strong effect on the outcome of cultural evolution.

For simplicity, we focus on two different kinds of learning: payoff-based learning (in which individuals tend to imitate individuals with higher payoffs) and frequency-based learning (in which individuals tend to imitate the majority of the group). In our models, we compare groups that are *homogeneous* with respect to social learning strategy (*i.e.*, all individuals employ the same mixed learning strategy that consists of both payoff-based learning and frequency-based learning) with groups that are *heterogeneous* in this respect (*i.e.*, each individual uses either only payoff-based learning or only frequency-based learning, but individuals differ with respect to which type of learning they use).

Both models have the same basic structure, closely following the setup of our primary experiment. We consider groups in which individuals are involved in social interactions. Individuals have a trait (A or B) that determines their behaviour in these interactions. Payoffs of A and B are $pa + (1-p)b$ and $pc + (1-p)d$, respectively, where p denotes the fraction of the group playing A . Individuals have pure strategies (either playing A or B), but their strategy can change over the course of time due to social learning.

S4.8. Details of the simulation model presented in the main text

For each of the four interaction contexts presented in Figure 4.4 of the main text, we simulated a process of cultural evolution in 10^5 replicate groups. In each of these replicates, individuals interacted in groups of eight and obtained payoffs according to the same payoff matrices as were used in the primary experiment (see section 2 of this Appendix). After each interaction, one individual was randomly drawn to update its strategy (see below). This cycle was repeated for 160 iterations (loosely according to the number of potential strategy updating events in the experiment, in which eight individuals could update their strategy in each of 20 rounds). At the beginning of each simulation, seven individuals played *A* and one played *B*. This setup of our simulation mimics the introduction of a deviant behaviour in a resident group of which all individuals have adopted the same behaviour. In the best choice context, this scenario mimics the introduction of a superior technology to a group. In the social dilemma, this scenario mimics the invasion of a defector in a group of cooperators. In a coordination and evasion games, this mimics the introduction of a deviant behaviour that, when rare, leads to lower and higher payoffs, respectively.

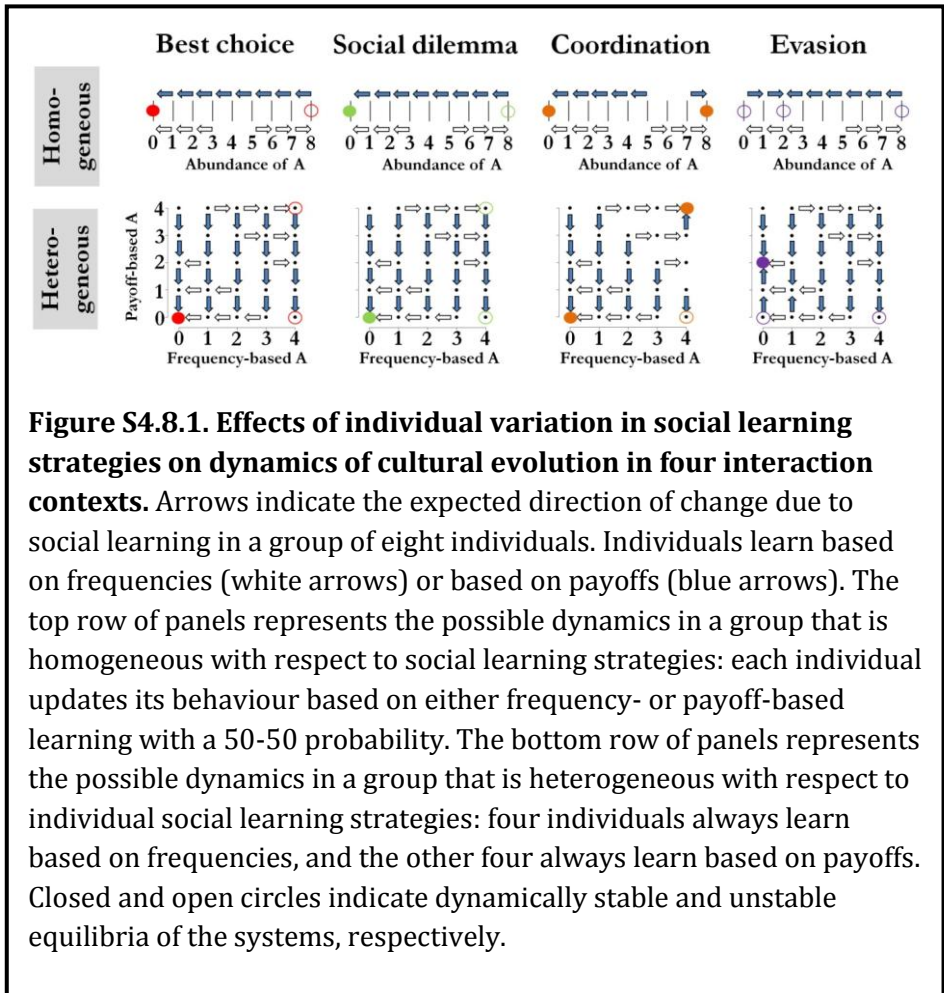
We implemented updating of traits through social learning in such a way that individuals used the same amount of information as was allowed in the primary experiment. In the event of frequency-based learning, behavioural updating was based on the previous decisions of a sample of six group members. The behaviour was updated to the behaviour exhibited by the majority of the sampled group members (if both behaviours were observed in equal proportions, the behaviour was not changed). For payoff-based learning, behavioural updating was based on both the previous decision and the previous payoff of a sample of three group members. The behaviour was updated to the behaviour that was associated with the highest payoff in this sample (but only if this payoff exceeded own payoff). In the simulation model, we do not allow for spontaneous switching between *A* and *B* (cf. mutation or innovation; see the alternative model in section S4.9 of this Appendix for a scenario in which we relax this assumption), neither do we add a stochastic component to the payoffs. These assumptions are made for simplicity. We documented the number of simulation runs that fixated for *A* and for *B*

(recall that there is no spontaneous strategy switching in our model, so a fixated group cannot change anymore).

In Figure S4.8.1, we illustrate how frequency- and payoff-based learning affect the expected change in the abundance of A and B in each of the four interaction contexts, in groups that are homogeneous and heterogeneous with respect to their social learning strategies. This graph aims to give insight into the dynamics that lead to the outcomes of cultural evolution presented in the panels in Figure 4.4 of the main text. In the best choice context and the social dilemma, the individually inferior behaviour A is more likely to fixate in a group in the absence of individual variation. The arrows in Figure S4.8.1 offer an intuition for this result: if a payoff-learner in the heterogeneous group has adopted the individually superior behaviour B , this individual will never switch back to A . This guarantees the eventual fixation of B . However, in heterogeneous groups, fixation in A always remains possible, particularly when B is initially rare and is disfavoured by frequency-based learning. In the coordination game, there are no differences between homogeneous and heterogeneous groups in terms of the outcome of cultural evolution. This is because any learning event (payoff-based learning or frequency-based learning) will lead to choosing A , so the fixation of A is the only possible outcome.

The most striking effects of individual variation in social learning strategies occur in the evasion game. The probability of fixation of either A or B is much higher in groups that are homogeneous. Note that in the evasion game, average payoffs are higher when A and B coexist (see section 2 of this Appendix). As illustrated by the bottom right panel of Figure S4.8.1, the dynamics of cultural evolution in the heterogeneous group will lead to the group composition where all frequency-based learners play B , whereas the payoff-based learners play A and B in equal proportions. When at this equilibrium, frequency-based learning events can no longer lead to the fixation of the most common strategy (B), because the only individuals that are playing A are payoff-based learners. In contrast, in a homogeneous group in which two players are playing A and six are playing B , fixation of B (through two consecutive frequency-based learning events) is still possible, and even likely to occur at some point. The emerging role differentiation in the heterogeneous group, in which all frequency-based learners play A and the payoff-based learners

play *A* or *B* with equal probability, ensures that the group retains behavioural polymorphism. In this particular interaction context, individual variation in social learning strategies increases average payoffs in a group.



S4.9. Details and results of an alternative analytical model

Here we present the details and results of an alternative model, in which we use an analytical approach to the same question: to what extent do individual differences in social learning strategies affect the outcome of cultural evolution? In this case, we focus only on the evasion game (for which we obtained the most interesting results in the simulation model). As in our simulation model, we compare groups that are homogeneous with respect to social learning strategy (all individuals use both frequency-based learning and payoff-based learning with equal probability) with groups that are heterogeneous in this respect (half of the group always uses payoff-based learning, the other half always uses frequency-based learning). Again, individuals exhibit either of two behaviours (A or B), and one randomly selected individual per time step can update this behaviour through social learning.

Let i denote the number of individuals that play A , with $0 \leq i \leq n$, where n denotes group size. For each state i we specify the probabilities of transition to states $i-1$ and $i+1$. These probabilities are entered in a transition matrix, the dominant right eigenvector of which gives its stationary distribution (Karlin and Taylor 1975; Otto and Day 2011), describing the long-run proportion of time that the group spends in each state.

In a homogeneous group, in which all individuals use the same mixture of frequency- and payoff-based learning, the probability that an individual that plays A switches to B is given by

$$\Pr(A \rightarrow B) = \frac{i}{n} \left[(1 - \varepsilon) \left(\frac{1}{2} C_i^- + \frac{1}{2} P_i^- \right) + \varepsilon \right], \quad (\text{S4.9.1a})$$

where $\frac{i}{n}$ is the probability that an A player is randomly selected from the group, and C_i^- and P_i^- denote the probabilities in state i of switching from A to B by frequency-based learning and payoff-based learning, respectively. In contrast to our simulation model, we allow for spontaneous strategy switching; the rate at which this occurs is given by ε .

Social learning occurs with the complementary probability $1 - \varepsilon$. Similarly, the probability that a B player switches to A is given by

$$\Pr(B \rightarrow A) = \frac{n-i}{n} \left[(1-\varepsilon) \left(\frac{1}{2} C_i^+ + \frac{1}{2} P_i^+ \right) + \varepsilon \right]. \quad (\text{S4.9.1b})$$

Social learning rules are deterministic: frequency-based learners always adopt the behaviour of the majority (when both types of behaviour are equally abundant, individuals stick with their original behaviour), payoff-based learners always adopt the behaviour that is associated with the highest payoff in the sample.

In a group that is heterogeneous with respect to the social learning strategies of its members, we cannot simply keep track of the number of A players in the group, since there are two distinct types of them (frequency-based learning A players and payoff-based learning A players). Therefore, we have to take four possible transitions from each state of the group into account: *i*) loss of a payoff-based learning A player, *ii*) loss of a frequency-based learning A player, *iii*) gain of a payoff-based learning A player, or *iv*) gain of a frequency-based learning A player.

Let i_c and i_p denote the number of frequency-based learning and payoff-based learners playing A , respectively. The total number of A players i in the group equals $i_c + i_p$. The probabilities of switching from A to B by either frequency- or payoff-based learning are given by

$$\Pr(A \rightarrow B)_c = \frac{i_c}{n} \left[(1-\varepsilon) C_i^- + \varepsilon \right] \quad (\text{S4.9.2a})$$

$$\text{and } \Pr(A \rightarrow B)_p = \frac{i_p}{n} \left[(1-\varepsilon) P_i^- + \varepsilon \right] \quad (\text{S4.9.2b})$$

respectively.

Similarly, the probabilities of switching from B to A by either frequency- or payoff-based learning are given by:

$$\Pr(B \rightarrow A)_C = \frac{n/2 - i_C}{n} \left[(1 - \varepsilon) C_i^+ + \varepsilon \right] \quad (\text{S4.9.3a})$$

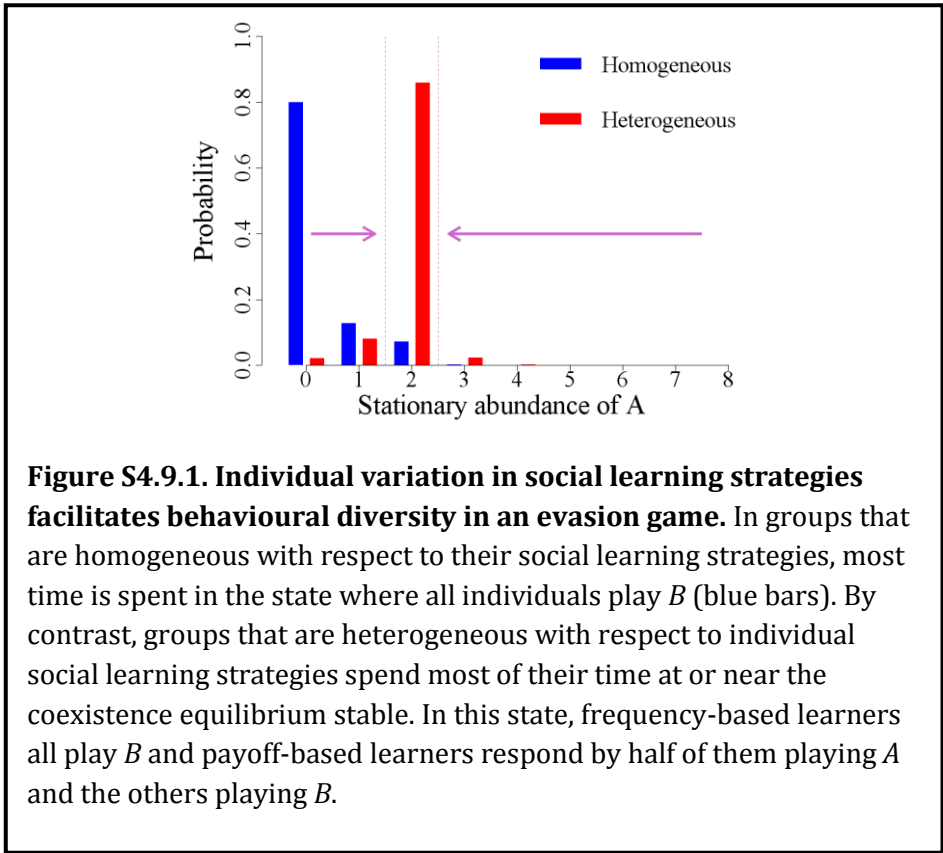
$$\text{and } \Pr(B \rightarrow A)_P = \frac{n/2 - i_P}{n} \left[(1 - \varepsilon) P_i^+ + \varepsilon \right] \quad (\text{S4.9.3b})$$

respectively.

Again, we enter those probabilities into a transition matrix, and calculate the dominant right eigenvector to obtain the stationary state of the group (Karlin and Taylor 1975).

We consider an evasion game with the same payoff matrix as the evasion game of our primary experiment in a group of eight players (see section 2 of this Appendix for details). The coexistence equilibrium of this game is located at $p^*=0.25$ (where two individuals play *A* and the other six play *B*). This means that payoff-learners that update their behaviour will adopt playing *A* or *B* if fewer or more than 2 individuals play *A*, respectively. They will stick to their previous behaviour if exactly 2 individuals play *A*. Frequency-based learners that update their behaviour will adopt playing *A* or *B* if the majority (more than four players) play *A* or *B*, respectively. They will stick to their previous behaviour if *A* and *B* are both played by four players.

Figure S4.9.1 shows the stationary states of the homogeneous and the heterogeneous groups. When there is no individual variation in social learning strategies, payoff-based learning will drive the group towards the coexistence equilibrium. However, frequency-based learning tends to drive the group towards the extreme state where all individuals play *B*. The net effect is that most of the time, all individuals play *B*. In contrast, as in the simulation model, a group with individual variation in social learning strategies allows a coexistence of *A* and *B* in the evasion game. At equilibrium, all frequency-based learners and two payoff-based learners play *B*, and the other two payoff-based learners play *A*. This confirms the findings of our simulation model, and supports the conclusion that individual variation in social learning strategies can affect the outcome of cultural evolution.



Part D. Experimental details

In this Part we describe the experimental procedures. It contains written instructions to the subjects, as well as instructions given on the computer screen during the experimental session. We conclude with screenshots of the decision making environment presented to the subjects.

S4.10. General experimental procedure

We ran 22 independent experimental sessions with $n=396$ subjects in total (primary experiment: eight sessions with sixteen subjects each; control: four sessions with sixteen subjects each; follow-up sessions: ten sessions with 20 subjects each). Experimental subjects were recruited by e-mail, using the subject pool of the Sociology department of the University of Groningen. Subjects (132 male, 264 female; aged 18 to 31,

mean age 22) were mostly undergraduate students from the social sciences, economics and biology. Experimental sessions lasted around 90 minutes, in which subjects earned €29.60 (£ 24.90) on average according to their performance, excluding a show-up fee of 5 euros.

Instructions were handed out in the reception room of the laboratory and read out loud by one of the experimenters. Each participant received a randomly chosen number corresponding to a desk in the laboratory. Before the experiment itself started, subjects played a test trial of five rounds, making them familiar with the decision making environment. After that, the four blocks of 20 rounds were started. At the beginning of each block of rounds, specific instructions for the upcoming interaction context were given on the computer screen, and participants had to fill out a brief quiz to check their understanding. Sessions finished with a questionnaire including items about personal demographic background. Participants were paid individually in the reception room.

S4.11. Instructions on paper

This section contains the instructions that participants received on paper. The on-screen instructions follow below, in section 12.

Welcome!

This session will last for approximately 2 hours. During the session **it is not allowed to talk** or communicate with the other participants. If you have a question, please raise your hand and one of us will come to you to answer it. During this session you will earn money. The amount you earn depends on your decisions and (sometimes) the decisions of others. At the end of the session the amount you have earned, plus a show-up fee of 5 euros, will be paid to you in cash. These payments are anonymous; you will be paid individually in a separate room. Please **stay seated at the end of the session** until your desk number is called. We will not inform any of the other participants about your earnings. It is impossible for us to associate your desk number with your identity.

Setup

In the coming session you own an imaginary farm. You will run this farm in **4 different experiments** in which you can earn points. At the end of the session, these points will be translated into real money (1000 points = 1 euro). At the beginning of each experiment, you will be grouped with 7 randomly chosen other participants in the room and each is randomly labeled with a number 1-8. All experiments are **anonymous**; you cannot find out the real identity of the other Farmers in your group, and they cannot find this out about you. During each experiment, your group consists of the same 8 participants, but when a new experiment starts, new groups are formed. At the beginning of each experiment, you will receive 2500 points to start with. The experiments are completely separate – what you earn in one experiment does not influence what you can earn in the other experiments. Also, it is likely that your earnings will vary between experiments.

Making decisions

Each experiment lasts for **20 Seasons**. In each Season, you will make a decision how to use the land on your farm. In each experiment, you will decide between two different options. All Farmers in your group will make this decision at the same time.

Once all Farmers have made their decision, all Crops of all Farmers in your group are sold on the Market. This procedure is automated in the computer program. The number of points you earn (your ‘revenue’) depends on three things:

1. **Crop value:** one of the two Crops is worth more points. This Crop yields higher revenues on the Market.
2. **Your land production:** the number of Crops you produce in a Season is affected by things that you don’t control, such as crop diseases or pests. In each Season, Farmers that make the same decision, will not receive exactly the same revenues. It will sometimes occur that a Farmer that chose the Crop with the higher value, will nevertheless get less revenues than a Farmer that chose the Crop with the lower value. However, on average, the

Crop with the higher value will give higher revenues. The variation in Crop production is fully determined by chance.

3. **Decisions of other Farmers:** in experiment 1, your revenues are independent of the decisions of the other Farmers in your group. In experiments 2, 3 and 4, your revenues **also** depend on the decisions of other Farmers. *How* your revenues depend on your decisions and the decisions of the other Farmers in your group, is given by the 'Market Rules'. These Rules are different in each of the experiments. The Market Rules will always be explained by the computer program at the beginning of each experiment. During an experiment, the Market Rules remain the same. Before the Seasons of an experiment start, you will fill out a short Quiz to check if you understand the Market Rules.

At the end of each Season, you are informed about the number of points you received for your Crops. It is also possible to lose points. The points you earn are stored in the computer memory.

Collecting information

Before making your decision, you can **collect information** about the decisions and revenues of the Farmers in your group. Three kinds of information are available about each Farmer:

- i) Decision in the previous Season
- ii) Revenue in the previous Season
- iii) Total points in the current experiment

At the start of each Season, you can indicate whether or not you want to collect information. Collecting information costs **25 points**. For those 25 points, you can collect up to 6 pieces of information in total. Remember that your revenues only *partly* depend on your land production in a Season. Collecting information about the decisions and revenues of other Farmers can be useful in finding out which Crops the other Farmers are producing, and which Crop has the higher value.

If you indicated not to collect information, no costs will be charged. In the first Season of each experiment, no information is available yet. At the start of the session, a **test trial** will make you familiar with the decision

making environment on the computer screen. This test trial lasts for 5 Seasons. You cannot earn anything; the points are not worth any money. After the test trial, the experiments will start.

Time limits

In each Season (except the first of each experiment), you have **10 seconds** to decide if you want to collect information about the Farmers in your group. If you have not decided within this time period, you will not get the opportunity to collect information. If you indicated 'Yes' (and confirmed), you have **20 seconds** to decide which pieces of information you wish to collect. In the Decision screen the information you collected is shown to you, and you can make your decision how to use your land. You have **20 seconds** to make this decision. If you have not decided (and confirmed your decision) within this time period, the computer program will make a random choice for you. When the Season has finished, a Results screen shows you the number of points you earned in this Season. This screen is shown to you for **10 seconds**. After the Results screen, a new Season will start.

After 20 Seasons, the experiment is over, and a new experiment will be started by us. At the beginning of this new experiment, you will again be grouped with 7 randomly chosen other participants in the room and each participant is randomly labeled with a number 1-8, and the new Market Rules are explained in the computer program. After 4 experiments, the session ends.

End of the session

At the end of the series of 4 experiments, you are asked to fill out a short Questionnaire. Once you have filled out this Questionnaire, you will be paid in a separate room. Please stay seated until we call your desk number.

S4.12. On-screen instructions

Before a new environment started, a screen was displayed to the subjects, announcing:

A new experiment starts now.

You are grouped with 7 randomly chosen other participants. You receive 2500 points to start with.

Subsequently, the 'Market rules' of the upcoming environment were displayed. There rules were specific of each of them, followed by a quiz to check the subject's understanding of the market rules. The quiz consisted of four statements, and subjects had to indicate if they were true or false. The rounds (or 'seasons') of the environment were started once all subjects had correctly filled out the quiz.

Best choice

Market rules

In this experiment, two Crops are available: **Potatoes** and **Wheat**.

One of the two Crops (Potatoes or Wheat) has a higher value than the other Crop.

The revenues of your Crops **do not** depend on the decisions of the other Farmers in your group.

Before the Seasons of this experiment start, a short Quiz will check if you understood the Market Rules.

Quiz

1. In each Season of this experiment, I can decide to produce Potatoes or Wheat.
2. Before I make my decision, I can collect information about the decisions of the other Farmers in my group.
3. The revenues of my Crop production also depend on the decisions of other Farmers in my group.
4. My group consists of the same Farmers during this experiment. When the experiment has finished, new groups will be formed randomly.

Social dilemma

Market rules

In this experiment, Apples are the only Crop available.

Each Season, you can decide if you want to **spray Pesticides** during the production of your Apples.

Using Pesticides will increase your land production.

However, the **more** Farmers decide to use Pesticides,

the **lower** the value of **all Apples** on the Market

(both Sprayed Apples and Unsprayed Apples).

Quiz

1. New groups have been formed randomly.
2. In each Season of this experiment, I can decide to use Pesticides to produce my Apples.
3. The value of all Apples is higher when I choose to use Pesticides.
4. The value of all Apples is highest, if all Farmers produce their Apples without using Pesticides.

Coordination

Market rules

In this experiment, two Crops are available: **Melons** and **Pumpkins**.

In principle, one of the two Crops (Melons or Pumpkins) has a higher value than the other Crop.

Yet, the Supplier of the Crop seeds will give more discount if more Farmers decide to produce the same Crop.

This means that the revenues of your Crops **also** depend on the decisions of the other Farmers in your group:

the **more** Farmers decide to produce a specific Crop, the **higher** its revenues.

Quiz

1. In this experiment, the revenues of my decision depend on the decisions of the other Farmers.
2. The more Farmers decide to produce a certain Crop, the more discount the Supplier will give. This will increase the revenues of this Crop.
3. The more Farmers decide to produce Melons, the higher the revenues of Melons will be.
4. The more Farmers decide to produce Pumpkins, the higher the revenues of Pumpkins will be.

Evasion

Market rules

In this experiment, two Crops are available: **Carrots** and **Onions**.

In principle, one of the two Crops (Carrots or Onions) has a higher value than the other Crop.

Yet, if a Crop is scarce on the Market, its value increases.

This means that the revenues of your Crops **also** depend on the decisions of the other Farmers in your group:

the **more** Farmers decide to produce a specific Crop, the **lower** its value.

Quiz

1. In this experiment, the revenues of my decision depend on the decisions of the other Farmers.
2. The more Farmers decide to produce a certain Crop, the lower its value.
3. The more Farmers decide to produce Carrots, the higher the revenues of Carrots will be.
4. The more Farmers decide to produce Onions, the higher the revenues of Onions will be.

Collecting information and making decisions

Before making their decisions in a Season, subjects could collect information from their fellow group mates. On the screen, subjects were asked the following:

Do you want to collect information before making your decision in this Season?

yes

no

If a subject ticked 'yes', a screen appeared in which he could collect information about his fellow group members (see below for screenshots). If subjects indicated 'no', they immediately went on to the decision screen. After all subjects made their decisions, a Results screen was shown, giving information about the payoffs acquired in the current Season. After 20 Seasons, subjects were also shown their total number of points collected in the current environment.

S4.13. Screenshots

Season 2 out of 20 Remaining Time 11

Collect up to 6 pieces of information! [Update information!](#)

Farmer 1	Farmer 2	Farmer 3	Farmer 4
Previous decision <input type="radio"/>	Previous decision <input checked="" type="radio"/>	Previous decision <input type="radio"/>	Previous decision <input type="radio"/>
Previous revenue <input type="radio"/>	Previous revenue <input checked="" type="radio"/>	Previous revenue <input type="radio"/>	Previous revenue <input type="radio"/>
Total points <input type="radio"/>	Total points <input type="radio"/>	Total points <input type="radio"/>	Total points <input type="radio"/>

Farmer 6	Farmer 7	Farmer 8
Previous decision <input checked="" type="radio"/>	Previous decision <input type="radio"/>	Previous decision <input type="radio"/>
Previous revenue <input type="radio"/>	Previous revenue <input type="radio"/>	Previous revenue <input type="radio"/>
Total points <input type="radio"/>	Total points <input type="radio"/>	Total points <input type="radio"/>

You are Farmer 5

Previous decision
 Previous revenue
 Total points

Figure S4.13.1. Collecting information of fellow group members. Subjects could check radio buttons to collect up to six pieces information of each other Farmer in their group separately. After checking the desired information, they could click the button 'Update information!'.

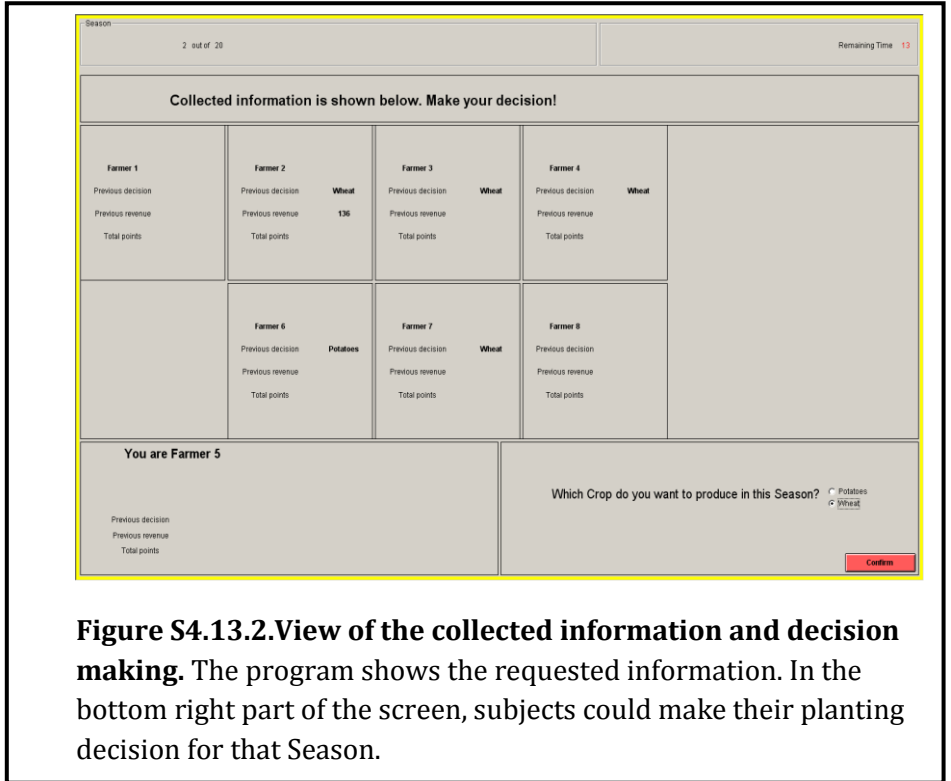


Figure S4.13.2. View of the collected information and decision making. The program shows the requested information. In the bottom right part of the screen, subjects could make their planting decision for that Season.

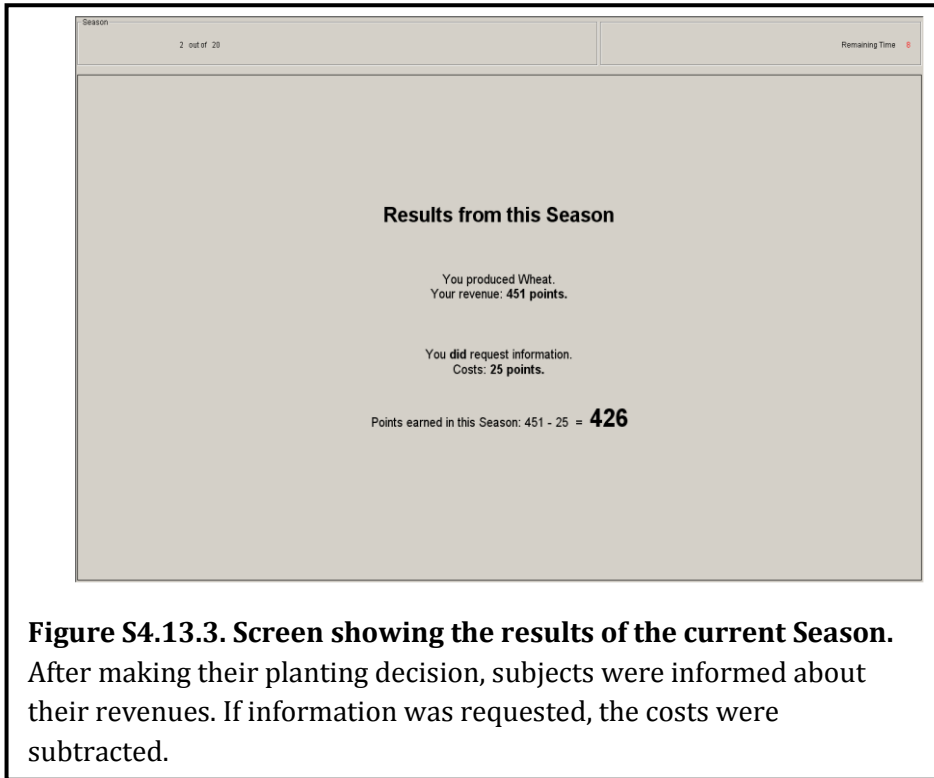


Figure S4.13.3. Screen showing the results of the current Season. After making their planting decision, subjects were informed about their revenues. If information was requested, the costs were subtracted.

Chapter five

Time consistency in human social learning strategies and the effects of their assortment on cooperative interactions

Lucas Molleman
Pieter van den Berg
Franz J. Weissing

Manuscript in preparation

Abstract

Social learning has caused humans to build up extensive cultural repertoires, allowing them to adapt to a wide variety of environmental and social conditions. Recent experimental evidence (reported in Chapter 4 of this thesis) shows that people differ with respect to their social learning strategies: some focus on the success of their peers, whereas others base their decisions on the behaviour of the majority. Moreover, this individual variation in social information use tends to be consistent across interaction contexts. However, the earlier findings leave open the question of whether these differences are stable over time. Here we provide experimental evidence that individual variation in social information use does indeed tend to persist over time, suggesting that social learning strategies are a structural part of individuals' behavioural repertoire. In addition, we show that the assortment of individuals with respect to their social learning strategies can affect the outcome of social interactions: groups of individuals that attend to majority behaviour (cf. using 'frequency-based learning') achieve high levels of cooperation in a social dilemma, as compared to groups whose members focus on the efficiency of peers' decisions (cf. using 'payoff-based learning'). In sum, our findings suggest that human social learning strategies are relatively stable, and that these strategies can have marked consequences for the outcome of cooperative interactions.

Introduction

Theories of cultural evolution assess how the frequencies and distributions of socially learned traits change over time (Cavalli-Sforza and Feldman 1981; Boyd and Richerson 1985). The dynamics of this process depends on how traits transmit between individuals. Theory suggests that different forms of social learning at the individual level (often referred to as ‘social learning strategies’; Laland 2004) can lead to different outcomes of cultural evolution at the population level (Chapters 2 and 3 of this thesis; Cavalli-Sforza and Feldman 1981; Boyd and Richerson 1985). For instance, it has been argued that groups whose members use conformist learning – preferentially copying the behaviour of the majority of one’s peers – are likely to sustain cooperation in a social dilemma, as compared to groups whose members use payoff-based learning (Boyd and Richerson 1985; but see Chapters 2 and 3 of this thesis).

The assumptions about human social learning strategies underlying theories of cultural evolution are largely circumstantial in nature, based on theoretical considerations and experiments from social psychology that were designed for other purposes. Only recently, systematic empirical evidence has started to accumulate on how people actually learn from each other in making decisions under uncertainty (McElreath et al. 2005; McElreath et al. 2008; Efferson et al. 2008; Mesoudi 2008; Mesoudi 2011). The results of these experiments suggest that people use both payoff-based and frequency-based social learning strategies in determining what to do. However, these studies typically focus on non-social contexts in which payoffs of behaviour do not depend on the behaviour of others (Traulsen et al. 2010 being the only exception known to us).

The experimental work presented in Chapter 4 of this thesis was the first to compare social learning strategies across social and non-social contexts. We observed that individuals differ strongly and consistently with respect to their social learning strategies. Some individuals tended to largely ignore social information in a range of different situations, whereas others frequently relied on social information in making their

decisions (Chapter 4 of this thesis, Figure 4.2a). We found even more pronounced differences in the type of information that individuals used: some consistently focused on the payoffs of their peers, whereas others tended to consistently disregard payoffs altogether, basing their behaviour on the frequencies of behaviour among their peers (Chapter 4 of this thesis, Figure 4.2b). These pronounced differences in social learning strategies persisted across contexts: individuals that tended to use payoff-based learning in one interaction context (*e.g.*, a social dilemma) were likely to use this form of learning in other interaction contexts (*e.g.*, a coordination or an evasion game; Chapter 4 of this thesis, Figure 4.3, S4.7.3).

The observation of distinct social learning strategies, and their consistency over interaction contexts raise two questions. First, do individual differences in social learning strategies persist through time? The observation of time consistent, distinct forms of social information use would suggest that social learning strategies can be viewed as ‘personality traits’, reflecting differences among individuals that stably persist across time and contexts. Second, how does assortment of distinct social learning types influence the outcome of social interactions? Human interactions are often structured in networks, and it seems plausible that, as a result of this structure, social learning strategies are distributed in a non-uniform way, so that individuals with similar social learning strategies are more likely to interact. The study presented here addresses these two questions by means of a set of decision making experiments.

Methods

Overview

We measured how individuals’ social learning strategies at two points in time, t_1 and t_2 , with four weeks in between. At t_1 , we documented each subject’s social information use in the four different interaction contexts (best choice, social dilemma, coordination game, evasion game). Subsequently, we categorized subjects based on reliance on payoffs in social learning, and assigned them to sessions for t_2 according to their social learning strategies. At t_2 , we again measured the subject’s social

information use in the four interaction contexts, using a slightly different setup. This procedure allowed us to *i*) quantify the time consistency of subjects' social learning strategies, and *ii*) assess the effects of assortment of social learning strategies on the outcome of group interactions.

Time point 1 (t_1)

The experimental sessions whose results are presented in the Appendix of Chapter 4 (Part B) of this thesis provide the basis for our analysis here. In these sessions at t_1 , a total of $n=200$ subjects were invited into the laboratory. In each of the ten sessions, twenty subjects participated and interacted in four groups of five. Subjects were confronted with four interaction contexts in sequence, in each of which the payoffs of decisions depended on the decisions of others in a different way (see below). In each interaction context, subjects made decisions in a block consisting of 20 rounds. At the beginning of each block, groups were randomly formed and instructions were given about the upcoming interaction context. In each round, subjects made a decision between two options (described as choosing to plant either of two crops on a virtual farm). Before making their decision, subjects could collect social information about their peer group members. From each of his or her four peers, a subject could collect information about *a*) the previous decision and *b*) previous payoff. The number of information pieces that an individual could collect before making a decision was limited to four. Each subject encountered, in random order, a sequence of *i*) a best choice situation, where one option yields on average a higher payoff; *ii*) a social dilemma, where one option increases the payoffs of all group members but in each situation yields lower payoffs for the individual; *iii*) a coordination game, where the payoff of each option increases with the number of peers choosing this option, and *iv*) an evasion game, where the relative payoff of either option decreases with the number of peers choosing this option. The payoff matrices of the interaction contexts are given in the Appendix of Chapter 4 of this thesis, section 2.

We documented for each subject in each interaction context his or her reliance on payoff information in social learning. To this end, we divided the number of information requests including payoff information over the total number of information requests. Based on these calculations, we

categorized each subject along the lines of the analysis used in Chapter 4 of this thesis: we distinguished ‘payoff-based learners’ (with a high reliance on payoffs in each of the four contexts), ‘frequency-based learners’ (with a low reliance on payoffs in each of the four contexts), and ‘sensible switchers’ (with a high reliance on payoffs in the best choice context, and a low reliance on payoffs in the other contexts). In case individuals could not be categorized according to any of the abovementioned types, they were labelled as ‘other’. We used this categorization to assign subjects to different sessions of t_2 , as will be explained below.

Time point 2 (t_2 ; one month later)

Based on the above categorization, we assigned payoff-based learners, frequency-based learners and sensible switchers to specific sessions of t_2 , planned four weeks after t_1 . Subjects labelled as ‘other’ were assigned to control sessions (for details of the assignment procedures, see Appendix, section 1). The participants in the control sessions were complemented by subjects unavailable at the time slots of their appropriate sessions. For each of the four categories we organized two sessions with 20 participants each ($n=40$ subjects for each category, $n=160$ in total). Subjects were not informed about the criteria we used to assign them to specific sessions.

The experimental setup of the sessions at t_2 was almost identical to the sessions of t_1 , the only change being the description of the decision making environment: instead of deciding to plant either of two crops on a farm, subjects had to choose between two investment options on a virtual stock market. This change aimed to minimize potential risks that individuals would fall back on a routine behaviour settled on at t_1 . Payoff matrices and the level of noise on payoffs remained unchanged.

In similarity to the sessions at t_1 , we sequentially confronted subjects with four interaction contexts in randomized order, each consisting of a block of 20 rounds. We again documented the types of social information that individuals requested before making their decisions. Comparing individual information use in each of the interaction contexts between t_1 and t_2 allowed us to assess the time consistency of social learning strategies. In addition, we tracked the group dynamics in each of the

interaction contexts, to assess the effect of assortment of types of social learning strategies.

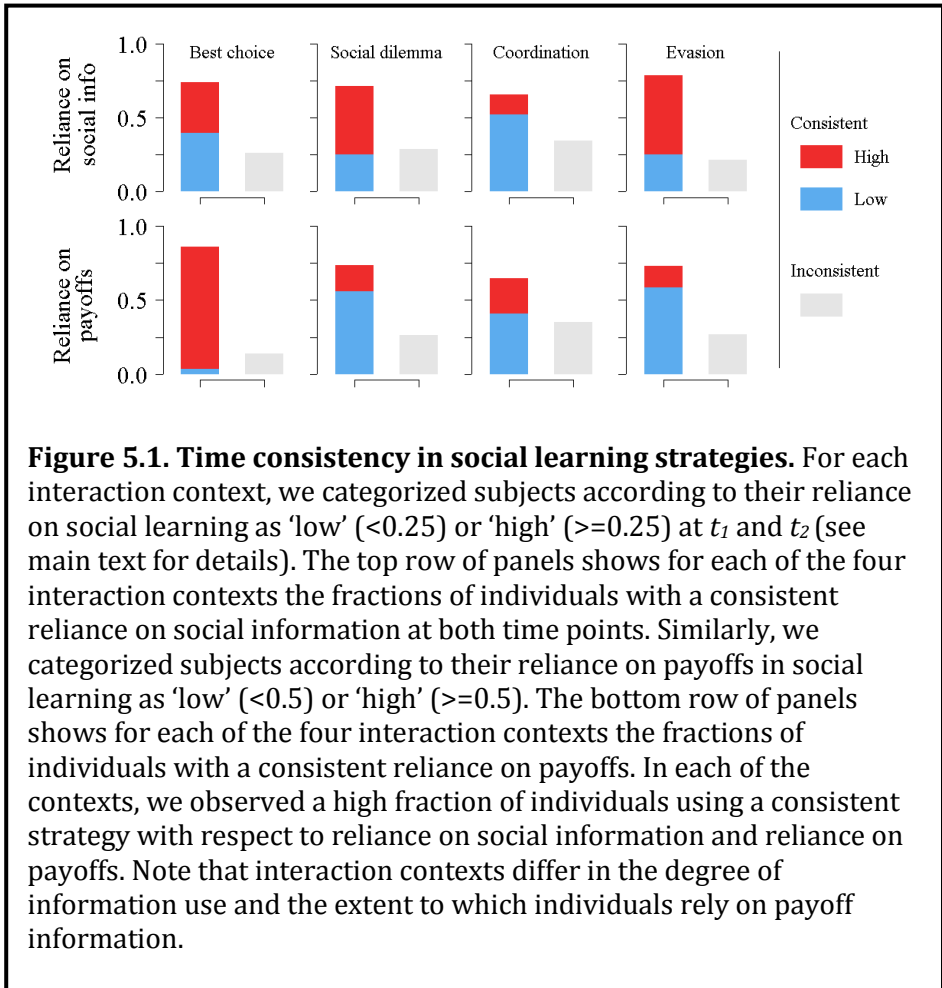
Results

Time consistency of social learning strategies

For both t_1 and t_2 , we categorized for each interaction context each subject's reliance on social information as 'high' (with a request rate of >0.25) or 'low' (with a request rate of <0.25). This 25% boundary splits up the population in roughly equal partitions for each of the interaction contexts (see Chapter 4 of this thesis, Figure S4.7.2a). The top row of panels in Figure 5.1 shows for each interaction context the fraction of individuals that are consistent in time in their reliance on social information. For each interaction context, we observe that the majority of individuals tends to be consistent with respect to their reliance on social learning (Figure 5.1 top row of panels; compare the coloured bars indicating the fraction of consistent individuals with the grey bars indicating the inconsistent individuals). A more fine-grained analysis based on linear regression revealed that regarding the degree of reliance on social learning, behaviour at t_1 is a fairly good predictor for behaviour at t_2 (with R^2 values for best choice: 0.39; social dilemma: 0.38; coordination game: 0.19; and evasion game: 0.42; see Appendix, section 2 for details). These results support the conclusion that individuals' reliance on social information tends to be stable over time.

The extent to which individuals relied on payoffs in social learning was analysed in a similar way: for each interaction context, each subject's rate of requests for payoffs was labelled as 'high' (when social information requests included payoff information at a rate ≥ 0.5) or 'low' (with requests including payoffs at a rate <0.5). This 50%-boundary forms an obvious threshold to distinguish individuals that predominantly focus on payoffs from others who ignore payoffs in more than half of their requests for social information, and it also splits up the population in roughly equal partitions in three out of four interaction contexts (see Chapter 4 of this thesis, Figure S4.7.2b). The bottom row of panels in Figure 5.1 shows for each interaction context the fraction of individuals with a consistent (high

or low) reliance on payoffs at t_1 and t_2 . Again, for each interaction context, we observe that the lion's share of individuals tends to be consistent with respect to their reliance on payoffs. In addition, we ran a separate analysis based on linear regression on the individual degrees of reliance on payoffs at t_1 and t_2 . This analysis revealed that for each context separately, a positive relationship exists between individual reliance on payoffs at t_1 and t_2 (with R^2 values for best choice: 0.13; social dilemma: 0.26; coordination game: 0.12; and evasion game: 0.19; see Appendix, section 1 for details). The results of both analyses suggest that individuals are also time consistent with respect to their reliance on payoff information.

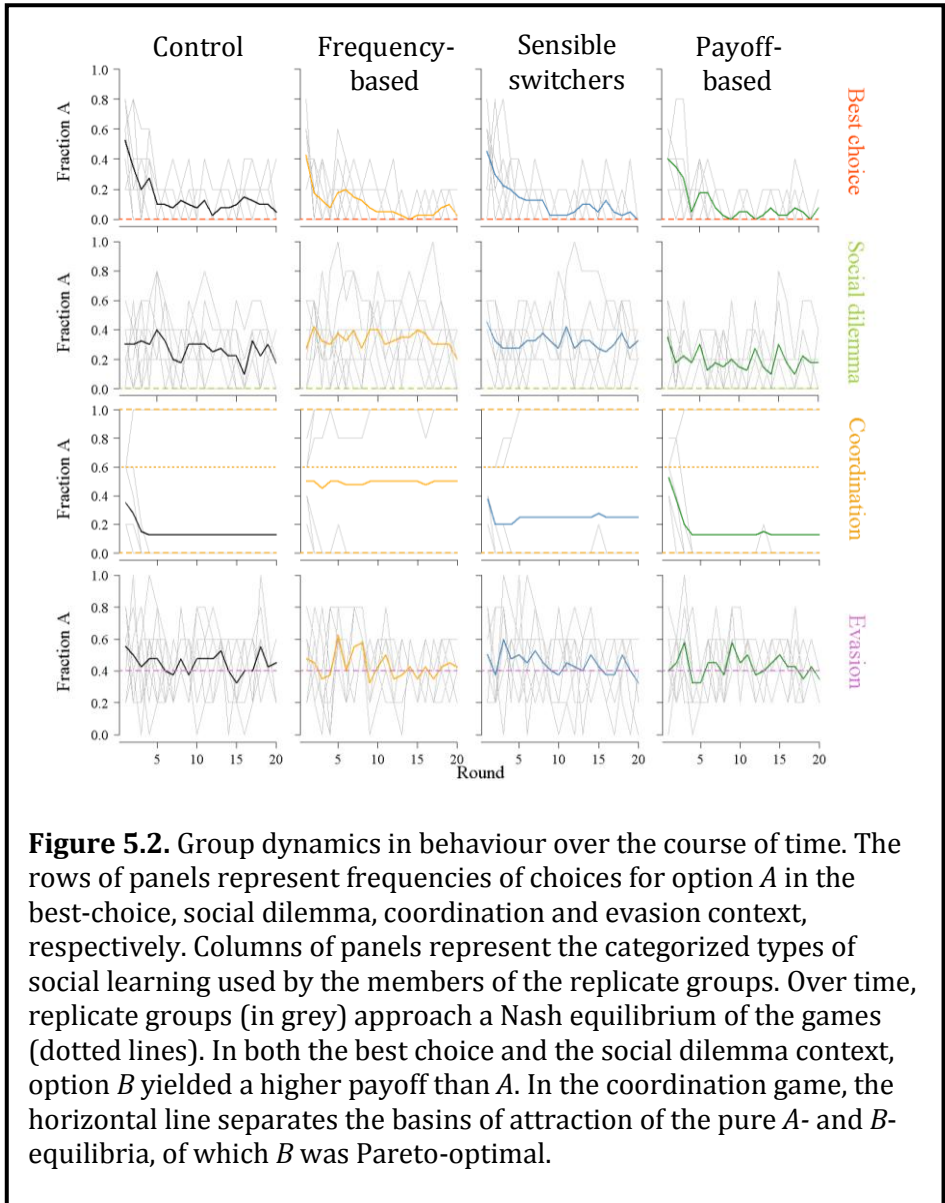


Information use differed between interaction contexts. For instance, reliance on social information tended to be high in the evasion game, as compared to the coordination game (Figure 5.1, top row; compare the relative sizes of the red and blue stacked bars in both games). This can be explained by the fact that in the coordination game, most replicate groups rapidly reached an equilibrium in which all members performed the same behaviour (see below decision dynamics). At this equilibrium, subjects had no incentive to switch behaviour, making (costly) social information less valuable. By contrast, in the evasion game, social information retained

its value as the frequencies of behaviour tended to fluctuate (see below, Figure 5.2, and Chapter 4, Figure S4.7.1). The extent to which individuals used payoff information in social learning also varied between interaction contexts: in the best choice situation, a majority of individuals consistently used payoff-based learning (bottom left panel of Figure 5.1; compare the relative sizes of the red and blue stacked bars); and this overall bias towards payoff information was more pronounced than we observed before (Chapter 4 of this thesis; Figure 4.2b, left panel). In the other three interaction contexts – where payoffs of decisions depended on the decisions of others – reliance on payoffs was lower considerably lower. This is in line with the findings reported in Chapter 4.

Effects of assortment of social learning strategies on the outcome of group interactions

Next, we turn to the effects of assortment of individuals with respect to these strategies on the outcome of social interactions within groups. We will focus on group dynamics in the sessions at t_2 only; for game dynamics in the sessions at t_1 , see Part B of the Appendix of Chapter 4 of this thesis. Figure 5.2 presents the average behaviour in each of the replicate groups. The rows of panels in this Figure show for each interaction context, for each category of individuals, average behaviour over the course of time.



Within each of the interaction contexts, groups tended to approach a Nash equilibrium of the one-shot versions of the game (Figure 5.2). For the best choice situation, the coordination game and the evasion game, we do not observe strong differences between assorted groups in terms of group dynamics. For the social dilemma, however, we observe a remarkable

difference in cooperation levels (highlighted in Figure 5.3). Groups whose members predominantly learn based on frequencies (in case of the social dilemma these are both the 'frequency-based learners' and the 'sensible switchers') achieved significantly higher levels of cooperation than groups of payoff-based learners (Tukey's test: $P=0.010$ and $P=0.022$ for comparisons of payoff-based learners with frequency-based learners and sensible switchers, respectively; logistic generalized linear mixed model fit to individual decisions, using 'learning type' and 'period' as fixed factors, and 'replicate group' as random factor). As an alternative approach to analyze the differences in cooperation levels in the social dilemma, we pooled together frequency-based learners and sensible switchers into one cohort of 'frequency-based learners'. Subsequently, we compared average cooperation rates in control sessions (8 observations), frequency-based learners (16 observations) and payoff-based learners (8 observations). This analysis confirms that mean cooperation levels in groups of frequency-based learners are significantly higher than those in groups of payoff-based learners (Tukey's test: $P=0.0125$). This result is in line with the earlier findings, and supports our conclusion that frequency-based learning can lead to high levels of cooperation, as compared to payoff-based learning.

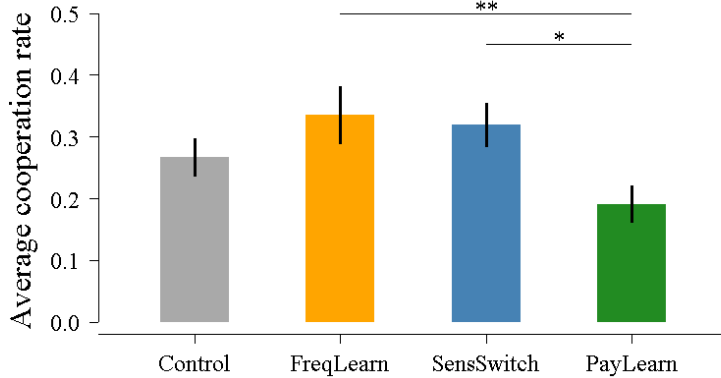


Figure 5.3. Mean cooperation rates in the social dilemma in each of the four treatments. Based on their information requests at t_1 , subjects were classified as ‘frequency-learners’ (largely ignoring payoff information of peers), ‘sensible switchers’ (only regarding payoffs in a best choice situation) or ‘payoff-learners’ (regarding payoff information in any context). Each bar shows averages of aggregated cooperation rates in the eight groups (± 1 SE; $n=8$ for each bar) in the experiment at t_2 . Groups that consisted of frequency-based learners or sensible switchers (orange and blue bars) achieved higher levels of cooperation than groups consisting of payoff-based learners (green bar). Asterisks refer to P -values of comparisons between groups (Tukey’s test: * $P<0.05$, ** $P<0.01$).

Discussion

Intriguingly, recent research has provided experimental evidence that subjects from different cultural backgrounds behave very differently when confronted with problems of cooperation (Henrich et al. 2001; Herrmann et al. 2008; Gächter and Herrmann 2009; Gächter et al. 2010). It is possible that distribution of other behaviours involving social interaction, such as social learning strategies, varies from culture to culture. For instance, it seems plausible that reliance on social learning may correlate with culture-level characteristics (Inglehart and Baker 2000). For instance, the degree of collectivism in a society may affect the degree to which individuals learn socially as opposed to non-socially, and may also influence conformist tendencies in deciding how to behave. Alternatively,

social learning strategies may depend on the degree to which people tend to generally trust others. Trust may affect the availability and reliability of useful social information: in low-trust societies, social information may be likely to be concealed or manipulated in a strategic way (*e.g.*, in attempts to gain competitive advantage), whereas in high-trust societies it might be easier to obtain reliable social information, allowing individuals to copy useful behaviours and to estimate the payoffs of one's peers. It is an empirical question whether social learning strategies indeed differ between societies, and correlate with one of the culture-level characteristics above. In case such society-level characteristics systematically influence social learning strategies of their members, societies likely differ with respect to the dynamics of cultural evolution (*e.g.*, how new technologies, beliefs, social norms, etc. spread).

Our experiments show that differences in social learning strategies occur on a smaller scale, within a single population. Our subject pool mainly consists of subjects from the University of Groningen, The Netherlands, aged 18-25. The coexistence of stable, distinct types within one demographically homogeneous subject pool is a phenomenon that calls for further theoretical and experimental investigation of the causes and consequences of this variation.

As a first step, we might ask how stable individual variation in social learning strategies within a population can emerge and be sustained in the face of natural selection. Models studying the evolutionary emergence and persistence of consistent variation in behavioural types could help explain the observed differences in social learning strategies (Wolf et al. 2007; Wolf et al. 2008; McNamara and Leimar 2010; Wolf et al. 2011; Wolf and Weissing 2012). Efficient behaviour in different interaction contexts may call for different forms of social learning. For instance, whereas focusing on payoffs seems useful in a best choice situation, it might be better to base decisions on the frequency of behaviours among your peers when the outcome of decisions depend on what others do. Adjusting their learning strategy from environment to environment could allow individuals to readily adopt optimal behaviour. However, there might be costs involved for being flexible (*e.g.*, in terms of the formation and maintenance of cognitive machinery, or time and deliberation costs of enhanced processing of (social) information). As a result, one might expect that

individuals are inclined to use some mixture of frequency-based and payoff-based learning. In our experiment, however, we observe pronounced differences between individuals with respect to the type of social information they use. It seems plausible that a consistent focus on one type of information (*i.e.*, payoffs or frequencies) could stem from a positive feedback mechanism reinforcing an individual's learning strategy. Over time, individuals may become more proficient in exploring their social environment and respond to specific types of social information in an adequate way. For instance, once individuals have gained experience in estimating the success of the behaviour of their peers, they may be inclined to exploit this potential advantage.

Time consistency in behaviour in cooperative contexts is an issue that only recently started to attract attention in the experimental social sciences. Evidence for individual stability of cooperative tendencies over time is mixed (Brosig et al. 2007; Volk et al. 2012). Our results provide empirical evidence for time consistency of individual social learning strategies. The fact that we did not find even greater degrees of consistency over time may be partially explained by the fact that our subjects were relatively young. Empirical studies have shown that in humans, consistency in personality traits increases as people grow older (Roberts and DelVecchio 2000). This suggests that we may find even higher levels of time consistency in social learning strategies in individuals from older age classes. Another factor that may influence degrees of consistency in social learning strategies is the time interval between observations (Roberts and DelVecchio 2000). The interval between t_1 and t_2 in our study was four weeks. This is relatively short compared with the total life span of an individual, but relatively long compared with more transient individual states that may affect behaviour in experiments, such as mood or fatigue. It remains an interesting question how stable social learning strategies – and traits affecting social behaviour in general – are over longer time spans, and which factors may cause them to change over an individual's life time.

Finally, assorting individuals with respect to their social learning strategies led to high cooperation rates in groups of frequency-based learners, as compared to groups of individuals who learn based on payoffs. Here, we offer a speculative explanation of this result. Focusing on

payoffs in social learning may reflect an individual's 'competitiveness'; *i.e.*, a subject's desire to perform better than others. Such a motivation is likely to be detrimental to cooperation in a social dilemma. Namely, whereas in repeated interactions cooperation may be beneficial when trust builds up and cooperation is reciprocated, defection is a way to secure the highest payoff relative to interaction partners. Results from a questionnaire associated with another cooperation experiment (Junikka et al., in preparation) support the suggestion that behaviour in social dilemmas may be in part motivated by competitive considerations. When asked to rate the importance of absolute versus relative performance in decision making, a substantial fraction of individuals rates relative performance as more important (see Box 5.1). Individuals that indicated that relative performance was an important factor in their decision making tended to contribute less to public goods. Together, these results suggest that at the individual level, links exist between human social learning strategies, strategic behaviour in social dilemmas and motivations in social interactions. Studies explicitly assessing these links seem a promising avenue of future investigation, which could add a new dimension to research on human personalities.

Appendix

This Appendix consists of 2 sections. In section 1, we give details of the procedure on which we based the group compositions in the sessions at t_2 . In section 2, we present additional analyses and measurements of individual time consistency in social learning strategies.

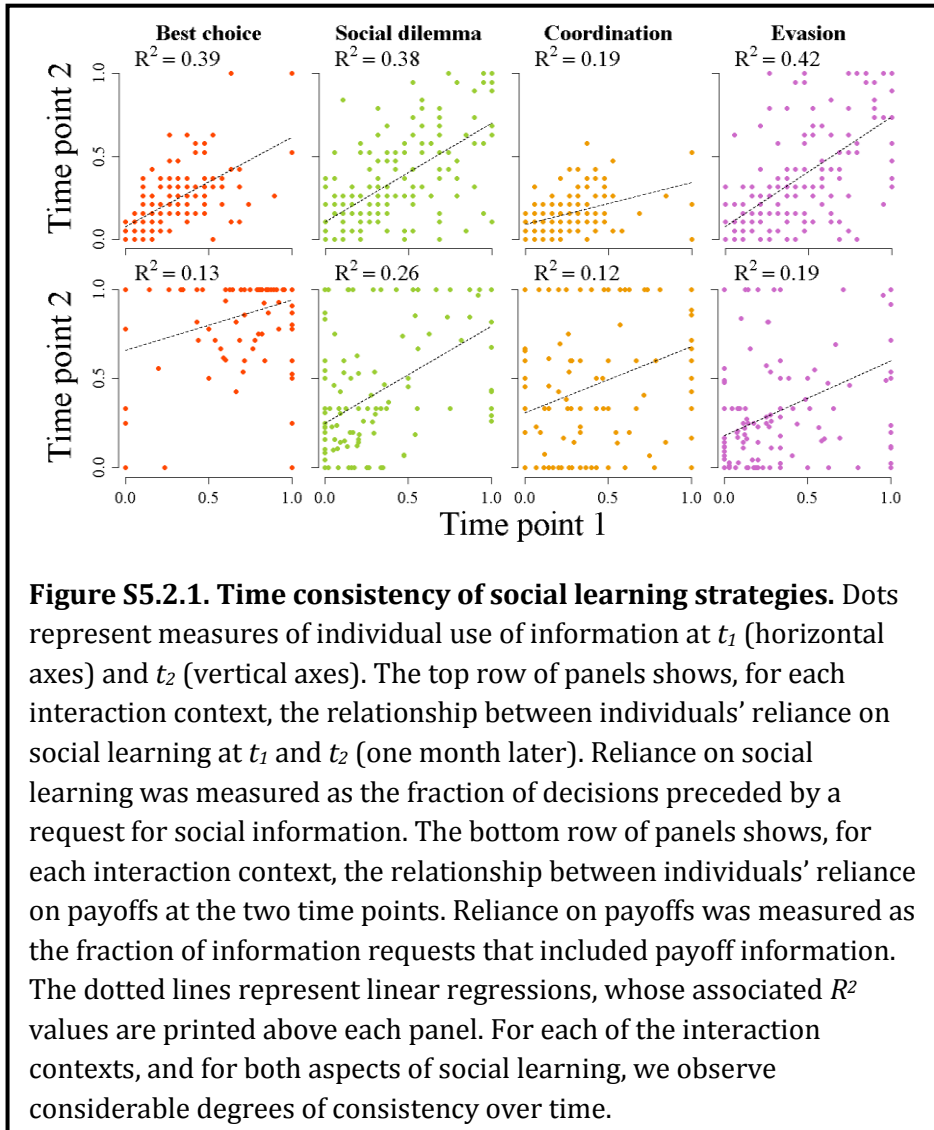
S5.1. Procedural details of assigning subjects to sessions of t_2

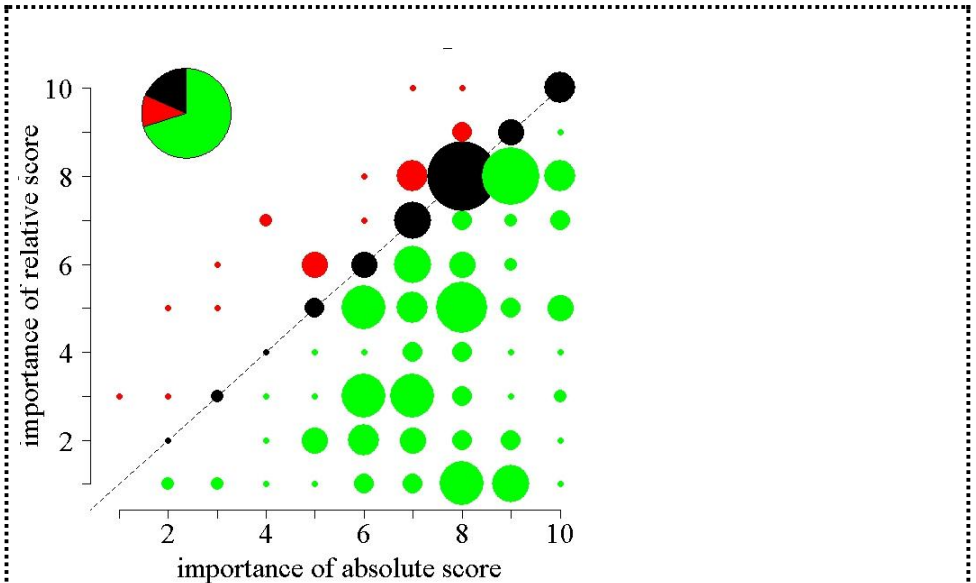
From the sessions of t_1 , we calculated for each subject his or her reliance on payoffs in each of the four interaction contexts. Subsequently, we assigned all individual to specific sessions of t_2 of the experiment, grouping together subjects with similar social learning strategies. The process of assigning categorized subjects to sessions was limited by two factors. First, at recruitment for the first phase, subjects were informed that they were expected to indicate their availability for sessions of t_2 for at least four – out of eight – predefined time slots (one month after t_1). To minimize constraints on assigning subjects to sessions, we rewarded each extra slot for which a participant indicated his or her availability with 1 euro. Second, despite the fact that social learning strategies were relatively extreme within each context, our measure of an individual's reliance on payoffs could take values anywhere between 0 (where a subject only requested information on behaviour, and ignored payoffs altogether) and 1 (where a subject always included payoffs in information requests). This led to the some degree of heterogeneity within categories. For instance, an individual categorized as a 'frequency-based learner' may have requested information about payoffs several times across the four interaction contexts. Despite this heterogeneity, the current procedure was strong enough to observe group-level effects of assortment in the social dilemma (Figure 5.3 of the main text).

S5.2. A more fine-grained comparison of social learning strategies at Time point 1 and Time point 2

In Figure 5.1 of the main text, we show that individual social learning strategies tend to be consistent over time: the fraction of consistent individuals is substantially higher compared to inconsistent individuals, for both reliance on social learning and reliance on payoffs. This analysis was based on a categorization of individuals' reliance on social learning and reliance on payoffs in each context (either 'high' or 'low' for both

aspects). Below we present an analysis on the uncategorized data, by and large confirming the findings presented in Figure 5.1. We find strong correlations in behaviour between t_1 and t_2 (Figure S5.2.1). This is the case for individuals' reliance on social learning, and, to a slightly lesser extent, for individuals' reliance on payoffs. This analysis supports our claim that social learning strategies tend to be consistent over time.





Box 5.1. Competitive motivations in a cooperation experiment. We conducted a decision making experiment on cooperation and conflict in which, in one of the treatments, subjects were confronted with ten rounds of a public goods game (Junikka et al. in preparation). In each round, groups of four were randomly formed, and subjects could choose to contribute some amount (0-20 monetary units) to a 'group project'. Points not contributed subjects kept for themselves. Summed contributions were doubled, and group members received an equal share of this doubled amount, irrespective of their contribution. This setup opposes individual and group interests: whereas the collective outcome is optimal when all group members contribute all their points, it is individually optimal to contribute nothing – thereby reaping the benefits of the contributions of others without paying the costs.

After the experiment, subjects filled out a questionnaire including the following two questions:

A. On a scale (1 – 10), how important was it for you to score a lot of points in this experiment?

and

B. On a scale (1 – 10) how important was it for you to score more points than other participants?

Box 5.1

Questions *A* and *B* aim to gauge the importance of, respectively, absolute and relative performance in decision making in a social dilemma. The plot shows the distribution of ratings of $n=240$ subjects. The size of the dots reflects the number of individuals with a specific combination of ratings on the two questions. The majority of individuals rates its absolute score higher than its performance relative to others (green dots). A vast majority fraction of subjects indicates that their relative performance was of some importance to them; only very few subjects answered question *B* with a '1'. Strikingly, a substantial fraction (>25%) indicated that their relative performance was at least as important as their absolute performance (the black and red dots; see the pie-diagram indicating the distributions of individuals over the respective categories). Moreover, some 10% of individuals indicated that their relative score was more important to them, as compared to their absolute score. A linear regression (taking the mean contributions in 10 rounds of a public goods game with random group formation as dependent variable, and questionnaire answers to Question *B* as the independent variable) reveals a strong negative relationship between the degree to which an individual was motivated by relative performance and its cooperation level in the public goods game ($\beta=-0.468$; adj. R-square = 0.104; $P<0.001$). This questionnaire result suggests that individuals are not exclusively driven by material self-interest; on the contrary, for a substantial fraction of individuals *relative* performance prevails over absolute performance.

Chapter six

Personal experience and reputation interact in human decisions to help reciprocally

Lucas Molleman
Eva van den Broek
Martijn Egas

Proceedings of the Royal Society B: Biological Sciences 280, 20123044

Abstract

There is ample evidence that human cooperative behaviour towards other individuals is often conditioned on information about previous interactions. This information derives both from personal experience (*cf.* direct reciprocity) and from experience of others, *i.e.* reputation (*cf.* indirect reciprocity). Direct and indirect reciprocity have been studied separately, but humans often have access to both types of information. Here we experimentally investigate information use in a repeated helping game. When acting as donor, subjects can condition their decisions to help recipients on both types of information at a small cost to access such information. We find that information from direct interactions weighs more heavily in decisions to help and participants tend to react less forgivingly to negative personal experience than to negative reputation. Moreover, effects of personal experience and reputation interact in decisions to help. If a recipient's reputation is positive, the personal experience of the donor has a weak effect on the decision to help, and vice versa. Yet, if the two types of information indicate conflicting signatures of helpfulness, most decisions to help follow personal experience. To understand the roles of direct and indirect reciprocity in human cooperation, they should be studied in concert, not in isolation.

Introduction

A key mechanism proposed to explain the evolution of cooperation is reciprocity (Fehr and Fischbacher 2003; Nowak 2006). When individuals interact repeatedly with the same partner, reciprocal donation can lead to mutual cooperation if the probability of future interaction is large enough (direct reciprocity; Trivers 1971; Axelrod and Hamilton 1981). This mechanism becomes less relevant as group size increases and thus the probability of two members meeting again decreases. In such situations, indirect reciprocity can favour cooperation if individuals can base the decision to help others on reputation, *i.e.* information about helping behaviour of individuals in previous interactions with others (Nowak and Sigmund 2005).

In humans, experimental work on helping behaviour confirms that both direct and indirect reciprocity can promote cooperation (*e.g.*, Wedekind and Milinski 1996; Wedekind and Milinski 2000; Seinen and Schram 2006). Such controlled experiments have focused on one of the two mechanisms in isolation, by employing designs that ensure the participant's information is either based on earlier interactions with the same individual (personal experience) or based on interactions of the recipient with other individuals (reputation). Other studies have compared effects of each of the two mechanisms, but in separate treatments. Some find higher helping rates towards recipients when information is available from personal experience rather than by reputation (Bolton et al. 2005). Others report the reverse (Dufwenberg et al. 2001). Simulations suggest that helping strategies based on reputation can survive in a population with other strategies based on personal experience (Roberts 2008). However, to the best of our knowledge, strategies based on both types of information were not yet taken into account. When making real-life decisions, humans typically use their own experience with another person and the experiences of others with this person. This raises the question how both types of information affect each other in the decision to help.

Here we report on an experiment in which subjects are allowed to condition helping behaviour on both personal experience and reputation. In a repeated helping game (Nowak and Sigmund 1998; Wedekind and

Milinski 2000; Seinen and Schram 2006; Engelmann and Fischbacher 2009) participants are randomly paired, and within pairs one participant is randomly assigned the role of donor and the other the role of recipient. In each game, the donor must decide to *i*) give a benefit to the recipient at a cost to himself or *ii*) to pass, resulting in no change in payoff for either of them. This helping game is iterated at least 100 times in groups of 12 subjects. To investigate the effects of personal experience and reputation, the game itself is preceded by an information stage in which donors can request costly information about their partner's past behaviour towards the donor (when roles were reversed) and/or towards others. Below, we refer to past behaviour towards the donor or others as 'direct' and 'indirect information', respectively. Donors choose to request either direct and/or indirect information at the same time, but may also decide not to request any information. A request for direct information provides the donor with the recipient's most recent decisions towards the donor, but not the order of decisions. For our baseline experiment, we impose 6 as a maximum to the number of decisions so that behaviour in the more distant past is ignored (*cf.* Seinen and Schram 2006; see Methods for details). Donors have no other way to retrieve this information, since subjects cannot be identified. A request for indirect information gives donors insight in the recipient's interaction history by presenting up to the last six decisions of this recipient when paired as a donor with other participants. Hence, indirect information excludes previous direct interactions with the donor. Additionally, we test for two factors that may affect the role of information in our baseline experiment. First, in reality, information on interaction histories with third parties is likely to be noisy due to mistakes in information transfer (*e.g.*, in gossip). To mimic this, we superimpose noise to the indirect information ('noise' treatment). Second, under random matching, indirect interactions are more frequent than direct interactions. To mimic this, we changed the available information on a recipient's history from six each to two direct decisions and twenty indirect decisions ('asymmetric' treatment).

To address the question of how personal experience and reputation affect helping behaviour we focus on the following three issues: Do humans base their decisions to help on personal experience as well as on reputations? How do personal experience and reputation affect helping behaviour? Are

the effects of personal experience and reputation on decisions to help independent?

Methods

We conducted a computerized experiment (programmed in Z-Tree; Fischbacher 2007) in laboratories at the CREED laboratory of the University of Amsterdam and the Sociology laboratory of the University of Groningen, the Netherlands. A total of 180 students were recruited to take part in 15 independent sessions with 12 participants each. In each session subjects were randomly assigned to cubicles. No communication was allowed. Written instructions for the experiment (in Dutch; see Appendix for an English translation) were provided and a quiz was used to ensure the subjects understood the instructions. Subjects knew that after 100 periods, a next period started with a probability of 0.9 (this was done to minimize end game effects). Each session lasted for approximately 90 minutes. In addition to a show-up fee of €7, subjects received an initial endowment of 3000 points (300 points = €1). Parameters were chosen such that requesting any type of information had a cost of five points, whereas helping had a cost of 150 to the donor, and a benefit of 250 to the recipient. To avoid income effects as much as possible, no information was given about the subjects' current earnings during the experiment. In principle, however, subjects could calculate their own earnings with pen and paper.

Whenever a participant was assigned the role of donor, two information request boxes were displayed. Choices to help or pass were shown as yellow and blue choices, respectively. When a donor requested information on previous decisions of the recipient, only the *number* of blue and yellow choices in the most recent decisions was given, not the order. In the baseline and the noise treatment, both direct and indirect information was limited to the six most recent decisions, giving subjects an opportunity to "clean their record" (*cf.* Seinen and Schram 2006). In our treatments we use random matching which implies for the baseline and the noise treatment that direct information is on average less recent than indirect information. Also, for direct information it would take a

larger number of rounds for an individual to clean its record compared to indirect information. The asymmetric treatment takes into account these arguments by limiting direct information to 2 decisions, and indirect information to 20 decisions, making both kinds of information equally recent.

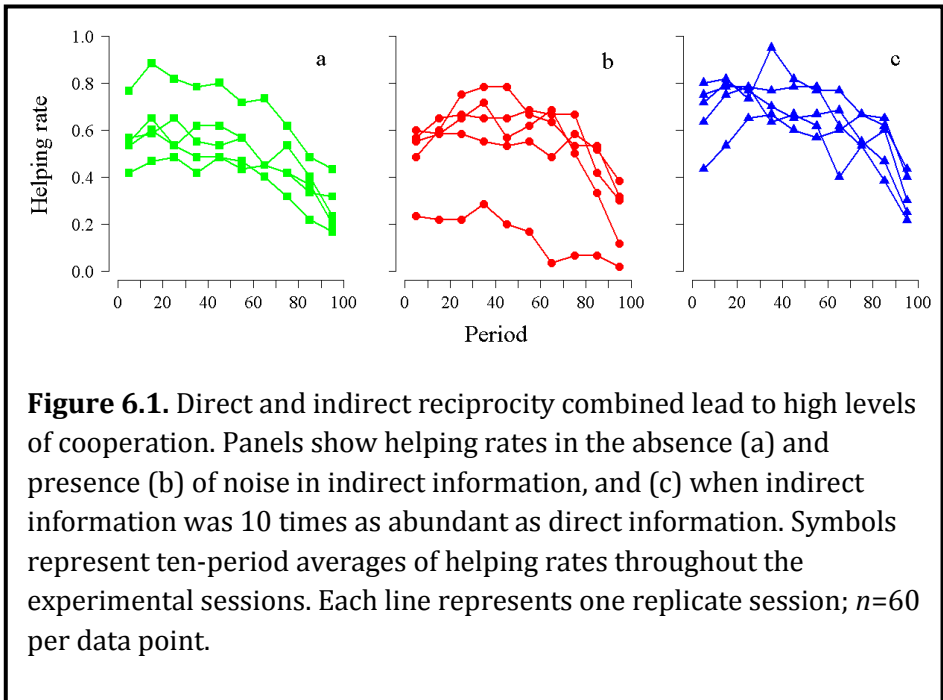
We impose a cost of five points for each information request, to ensure that subjects deliberately click on the information they are interested in. In the early phase of the experiment, the total number of decisions presented to the donor may be smaller than the maximum, (6 and 6 in the baseline and noise treatment; or 2 and 20 in the asymmetric treatment) or even zero. After decisions have been made, the donors and recipients are informed about their earnings in that period. At the end of the experiment, subjects completed a short questionnaire with topics such as age, gender and educational background. On average subjects earned €36.75.

Requested information of direct interactions was fully accurate. We added some noise to the indirect information in 5 of the 15 sessions to mimic distorting effects of gossip. One out of six pieces of information was randomly chosen and its content was flipped, *i.e.* from yellow to blue or *vice versa*. This implementation is slightly biased against extreme scores, since a 5:1 score is more likely to stem from 6:0 than from 4:2.

Statistical analyses were conducted in R (Venables and Smith 2005). We used a mixed-effects logit regression models to analyze how decisions to help or requests for direct and/or indirect information depended on various factors. We analysed these generalised linear mixed models (glmm) using R-package 'lme4' (Bates et al. 2012). All glmm analyses included controls for $\text{period}/100$, and $\text{period}^2/100^2$, correcting the model for temporal patterns in the experiment (*e.g.*, decreasing helping rates in the ending phase of the experiment). Unless stated otherwise, subject was included as random factor, nested in session replicate. As a check of the robustness of the effects obtained, we fitted another mixed-effects logit model to all data, including those decisions preceded by no information. In order to test for the robustness of the statistical findings, we verified that found effects are insensitive to the assumed cumulative probability distributions to help (*i.e.* probit vs logit).

Results

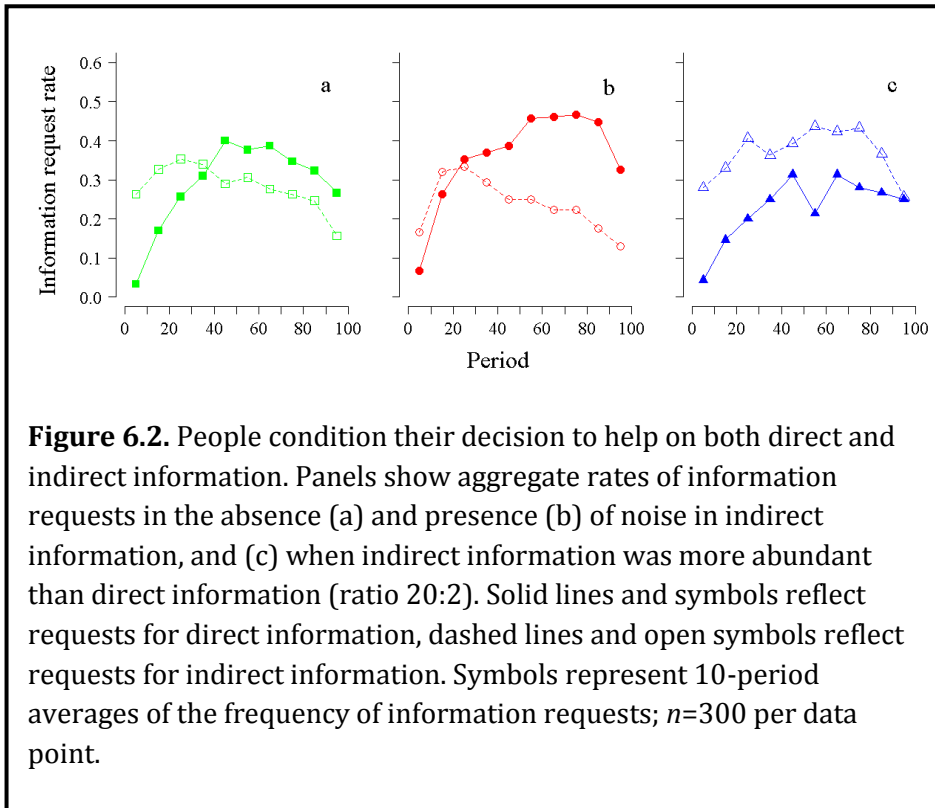
Subjects displayed a high degree of helpful behaviour in all treatments of our experiments (Figure 6.1). Helping rates did not differ between treatments (Table S6.1a), and in all replicate sessions of the three treatments helping rates declined towards the end of the experiment (Table S6.1a).



Subjects use information on both direct and indirect interactions

Averaged over all sessions, donors based 49.3% of their decisions on past decisions of the recipient. In an early phase, information on interaction histories is likely to be scarce and is therefore less frequently requested (Figure 6.2; Table S6.1b,c). This is particularly true for direct information. Initially, under random matching, subjects are unlikely to have interacted with each other before; as the session progresses, direct information becomes relatively more useful and is therefore requested more often (Table S6.1). The overall proportion of requests for indirect information

tends to decrease over time (Table S6.1d). Overall, adding noise to indirect information decreases its relative proportion of requests compared to direct information (Figure 6.2b, Table 6.2d). When indirect information is available in higher amounts, it is requested relatively more often (Figure 6.2c; Table S6.1d). When we control for variation at the individual level, treatments do not significantly differ in request rates for direct information (Table S6.1b). However, subjects do tend to increase requests for indirect information in the asymmetric treatment (Figure 6.2c, Table S6.1c). In the Appendix we describe patterns of variation at the individual level in the use of information, helping rates and performance in the experiment (Figure S6.1-S6.5).

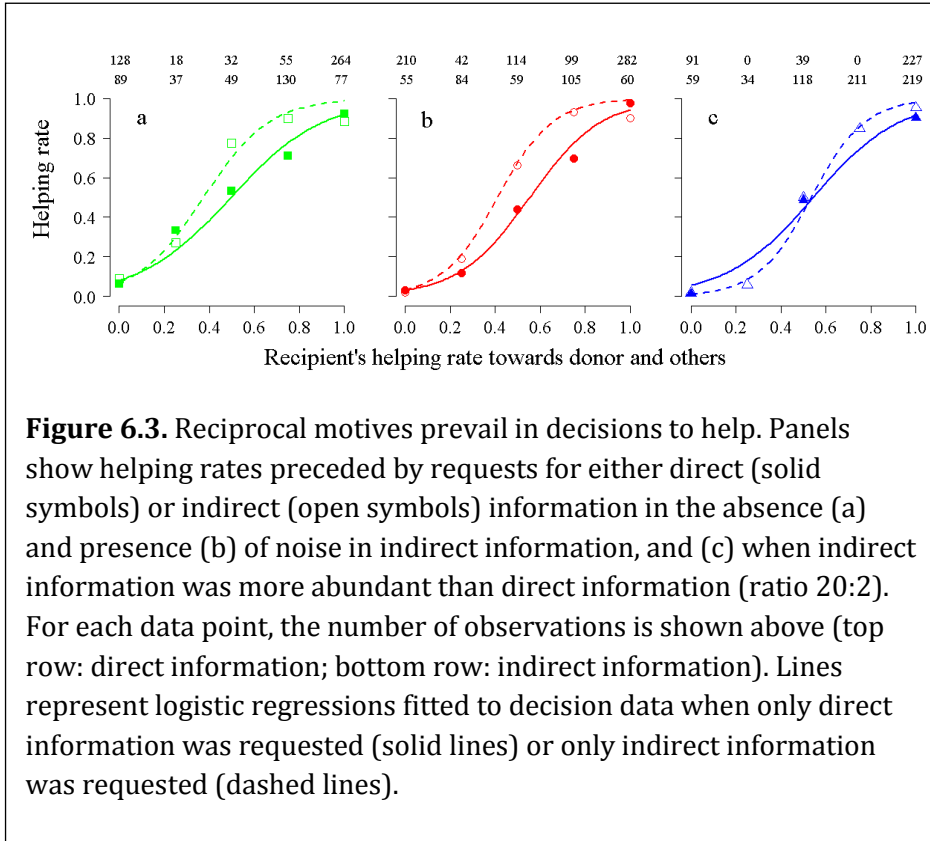


Direct and indirect information affect helping differently

For both direct and indirect information, we find a strong positive correlation between the recipient's helping rate and the rate of being helped (Figure 6.3, Table S6.2). The two types of information differ, however, in their effect on helping behaviour. We illustrate this with an analysis that contrasts the decisions preceded by requests for direct information only with the decisions preceded by requests for indirect information only. A GLMM fit to data detects a significant interaction effect between the observed helping rate of the recipient and the type of requested information (direct or indirect), indicating that direct and indirect information affect decisions to help differently (Table S6.2a). When we consider the difference between the number of helping and passing decisions (instead of helping rates), we find the same effects (Table S6.2b). On the same range of 'image scores', direct information has a steeper slope in all treatments (Figure S6.6), suggesting that subjects react to direct information in a more extreme way. No relationship was found between individual's reliance on reputation and overall helping rate (Figure S6.4).

Direct and indirect information are not independent

When we consider decisions preceded by requests for both direct and indirect information, we find that reciprocity effects are significant: positive information (helping the donor in the past, helping others in the past) has positive effects on helping rates (Table 6.1). The estimated effect of direct information ('help donor') is larger than the estimated effect of indirect information ('help others'), confirming the findings in section 2. Also, whenever a recipient's helping rate from direct and indirect information show values on opposing extremes, subjects follow direct information more often (number of helping versus passing decisions equals 41:18 when direct information is positive (>0.5) and indirect information is negative (<0.5). In the reverse case, helping versus passing equals 24:29; Fisher exact test: $p=0.013$).



The statistical model shows a significant interaction effect between direct and indirect information. The estimate of the interaction effect is negative, i.e., higher observed helping rates in either type of information are associated with a lower impact of the other type on helping behaviour. The effect of indirect information becomes weaker as direct information is more positive (and vice versa; Table 6.1). When direct information about the recipient indicates this subject often helped the donor, indirect information has a smaller effect on the decision of the donor to help than when the recipient hardly helped the donor. Similarly, when indirect information reveals the recipient often helped others, the effect of direct information is mitigated (Figure 6.4, Table S6.3).

	Estimate	<i>P</i>
Intercept	1.130	0.266
Period	3.842	0.223
Period ²	-5.561	0.033
Help donor	4.620	<0.001
Help others	3.838	<0.001
Help donor x help others	-3.883	0.006
Noise in indirect information	-0.193	0.793
Asymmetric information	-0.605	0.401

Table 6.1. Direct and indirect information are not independent. Values represent regression model estimates of factors affecting decisions to help. The analysis only considers decisions preceded by requests for both direct and indirect information. Significant effects include the recipients' fraction of helpful decisions towards the donor and others, as well as their interaction. Treatments do not affect helping rates significantly. Subject nested in session was included as random effect. Helping rates are centred before entering analysis. Period (squared) was included to control for the decreasing trend in helping rates over time.

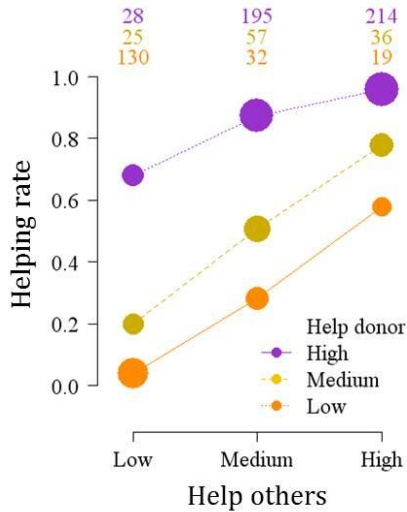


Figure 6.4. Personal experience and reputations interact in decisions to help. Data points correspond to rates of receiving help when both direct and indirect information were requested, for various categories of recipients' helping rate. The recipients' helping rate towards the donor (Help Donor) and towards third parties (Help others) are both categorized as 'Low' (≤ 0.3), 'Medium', (0.3-0.7), or 'High' ($>= 0.7$). The sizes of the plotted dots indicate the natural log of the number of constituent data points. These numbers are also given above the graph. Receiving rates are plotted as a function of helping rates towards others, for the three categories of Help Donor. Effects of indirect information are largest for low values of Help Donor (orange line). When direct information provides a high signature of helpfulness (purple line), indirect information has weaker effects (compare the slopes of purple and orange lines). Thus, the effect of reputation is weak when personal experience is good. A similar pattern emerges in the reverse case: when helping rates towards others are low, the effect of direct information is stronger (data points on the left hand side, where 'help others' is low, are closer to each other than the data points on the right hand side, where 'help others' is high).

Discussion

Our results show how people integrate direct and indirect information about past social interactions (Figure 6.2), confirming the commonly held belief that people use both personal experience and reputations to decide whether or not to help others. Even though information is costly, subjects often condition their decisions whether or not to help a recipient on the decisions that this person made previously. Cooperation levels are far above zero in all sessions (Figure 6.1), which is in agreement with a comparable experiment in which subjects always received information but could not distinguish direct from indirect information (Seinen and Schram 2006). The past behaviour of recipients, reflected in direct and indirect information, has strong effects on decisions of donors to help (Figure 6.3, S6.6, Table 6.1, S6.1a, S6.2, S6.3), marking direct and indirect reciprocal motives. It has been argued that indirect reciprocity is likely to wane in large groups because of the noisy nature of information, *e.g.*, due to gossip (Engelmann and Fischbacher 2009, but see Sommerfeld et al. 2007, 2008). By contrast, we find that people show sensitivity to noise not by a lower propensity to help but, at the aggregate level, people react to noise by higher proportions of requests for direct information (Figure 6.2, Table S6.1d). In the asymmetric treatment subjects rely more on indirect information (Table S6.1d). Despite the fact that under these conditions direct information was less information-rich (ratio 2:20) but equally costly, subjects frequently rely on direct information in making their decisions to help. This suggests that people have a preference to rely on personal experience, even when it gives a limited picture of a recipient's helpfulness.

Subjects react more strongly to information from personal experience than to reputational information (Figure 6.3a, b, S6.6, Table 6.1, Table S6.2, S6.4). This effect does not appear to be due to differences in individual strategies; individuals relying to different degrees on direct and/or indirect information did not show differences in helping rates (Figure S6.4). The inherent contrast between the direct and concrete effects of first-person experience and the indirect and abstract notion of reputation, however, may partly explain the observed pattern. Refusing to help may infuriate a direct partner, whereas the same thing happening to others does little more than eliciting empathy. Alternatively, the relative

impact of direct and indirect information could depend on the type of interactions that subjects encounter in their everyday life. For example, public information may guide behaviour of individuals from a small village. In a large city, however, interaction partners are mostly reputation-free strangers to each other. Hence, our results may be influenced by the fact that most participants were students from large cities.

Our results show that reputation has the potential to substantially increase helping rates when personal experience is negative (Figure 6.4). Conversely, however, when personal experience is positive, a bad reputation does not lead to severe decreases in helping. These results are robust to changes in assumptions on the evaluation of (in)direct information in the regression models (see Appendix). When personal experience is positive, helping rates are generally high, and a good reputation has little scope to increase helping rates even more. Such ‘ceiling effects’ may partly explain the detected interaction between personal experience and reputations. We believe that these effects are not particular to our experiment. In real life personal experience and reputations may compensate for each other: people might help others with a bad reputation if personal experience is good. Conversely, people may be more helpful to others after a bad personal experience, if these others have a good reputation (*cf* Figure 6.4).

In our experiment, as in previous experiments on reciprocity (*e.g.*, Wedekind and Milinski 1996; Wedekind and Milinski 2000; Seinen and Schram 2006), the aim is to measure effects of first-order information. Hence, subjects do not know whether previous decisions of their interaction partners are based on any information, and it is impossible to examine if a donor reciprocated the earlier decisions of the recipient. In our experiment it was thereby impossible to employ (theoretically superior) strategies that make use of second-order information, such as the reputations of subjects that were refused help by a recipient (*e.g.*, the ‘standing’ strategy; Leimar and Hammerstein 2001; Ohtsuki and Iwasa 2004; Ule et al. 2009). It seems plausible that insight in second-order information will influence helping behaviour: subjects might forgive a refusal to help more easily when they themselves refused help to a recipient before (but see Milinski et al. 2001). In addition, second-order

information is likely to affect the frequency and distribution of information requests, as well as the payoffs of reciprocal strategies. Our experimental design can be readily extended to address such scenarios.

Previous theoretical and experimental work on mechanisms driving human cooperative behaviour has focused on direct and indirect reciprocity in isolation. Building theory on the basis of these two mechanisms would require assumptions on how the two types of information, direct experience and reputation, are integrated in social decision making. However, how humans integrate these pieces of information is unknown, as yet. Our experiment shows that the effects of information on personal experience and reputations are interdependent in determining whether or not to help others: cooperation prevails when either personal experience or reputation is positive. This suggests that a higher level of cooperation can be attained, relative to a situation where only one of the two types of information is available. To understand the roles of direct and indirect reciprocity in human cooperation, these mechanisms should be studied in concert, not in isolation.

Appendix

Regression models fit to data underpinning claims in the main text

Table S6.1. Results of four regression models fit to all data, supporting the claims in the results section of the main text. Columns *a-c* show effect estimates of generalised linear mixed logit models (with associated *p*-values in brackets) using subject nested in session as random effect. The factors including ‘period’ account for temporal patterns in the data. The factors ‘noise’ and ‘asymmetric’ refer to the two treatments that were run in addition to the baseline treatment. All models detect strong temporal patterns. Furthermore, the asymmetric treatment shows significantly higher levels of requests for indirect information. Column *d* shows effect estimates of a generalised linear logit model that tests the relative use of direct and indirect information between treatments (cf. Figure 6.2 of the main text). The model is fit to the overall per-round number of requests for indirect information divided by the total number of requests for information, aggregated over all sessions per treatment. As suggested by Figure 6.2 of the main text, at the aggregate level, adding noise to indirect information leads to lower overall request rates for this information. The asymmetric treatment shows a higher overall rate of requests for indirect information.

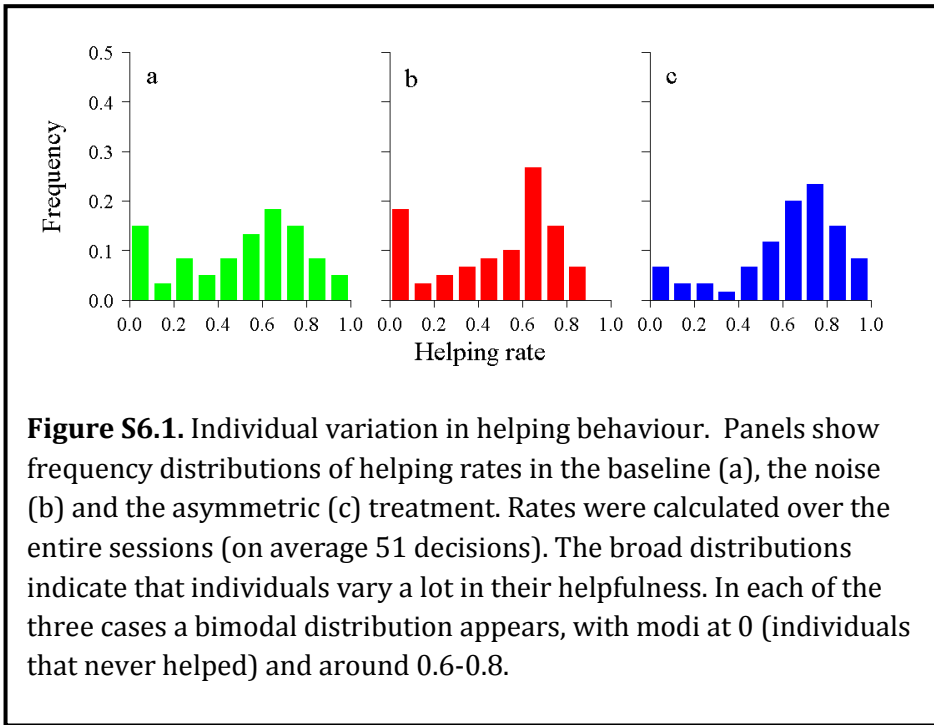
	(a) Decision to help	(b) Request for direct information	(c) Request for indirect information	(d) Proportion of requests for indirect information
intercept	0.215 (0.598)	-6.118 (<0.001)	-2.362 (<0.001)	1.205 (<0.001)
period / 100	2.439 (<0.001)	14.607 (<0.001)	4.105 (<0.001)	-4.060 (<0.001)
period^2 / 100^2	-4.760 (<0.001)	-11.912 (<0.001)	-5.014 (<0.001)	2.574 (<0.001)
noise	-0.223 (0.693)	0.736 (0.280)	-0.505 (0.322)	-0.395 (<0.001)
asymmetric	0.898 (0.112)	-0.508 (0.460)	1.016 (0.043)	0.555 (<0.001)
n =	9180	9180	9180	1496

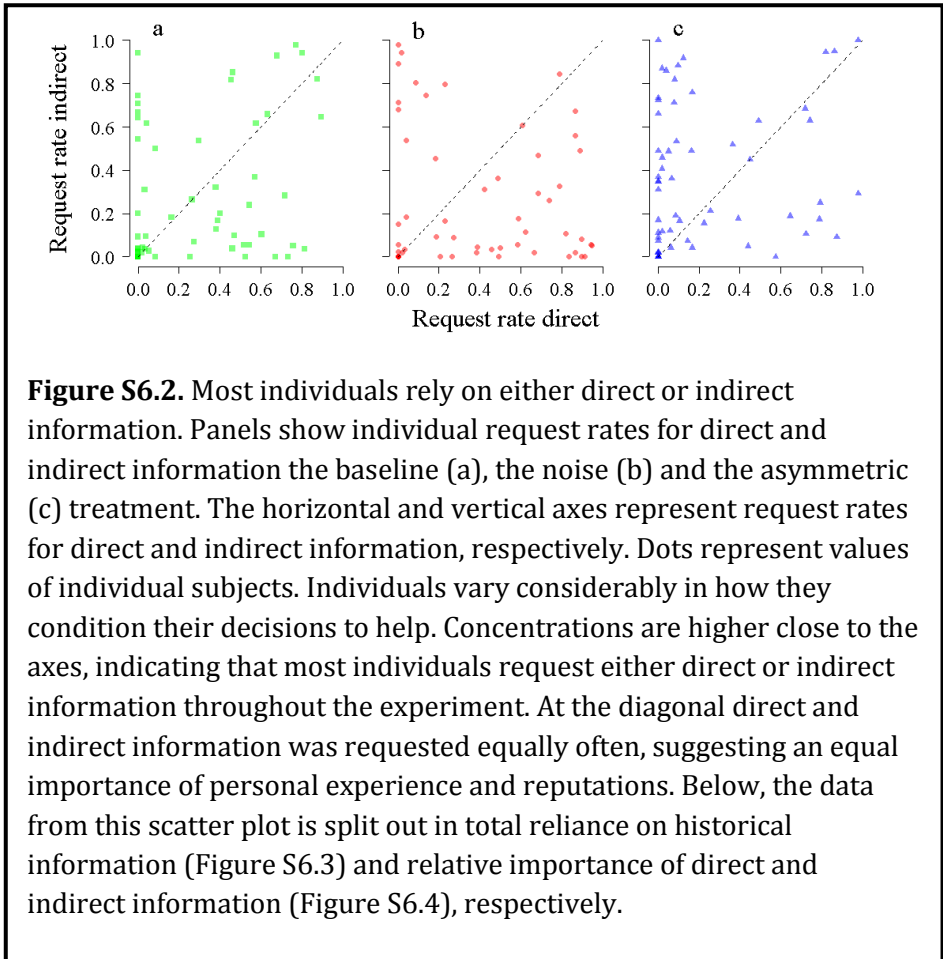
Table S6.2. Direct and indirect information affect decisions help differently. Columns present effect estimates of generalised linear mixed logit models fit to decisions that were preceded by requests for either direct or indirect information. Subject nested in session was included as random effect. Positive information on the recipient's helpfulness leads to higher levels of helping (*cf* a strong positive estimate for 'help donor or help others'). We observe a significant interaction effect between the observed helping rate of the recipient and the type of information, indicating that subjects react differently to personal experience than to reputations. The same effects are obtained when we consider not the fraction of helping (as opposed to passing) decisions, but the difference between the two (the number of helping minus passing decisions 'image scoring', column *b*).

	(a) Decision to help (fractions)	(b) Decision to help (image scoring)
Intercept	-3.540 (<0.001)	0.697 (0.010)
Period / 100	3.213 (0.021)	0.915 (0.248)
Period² / 100²	-4.414 (<0.001)	-2.913 (<0.001)
Noise	-0.177 (0.597)	0.012 (0.962)
Asymmetric	-0.584 (0.089)	-0.265 (0.307)
Help donor or help others	7.238 (<0.001)	5.809 (<0.001)
Info type (direct or indirect)	-0.577 (0.132)	0.039 (0.811)
fraction help x info type	1.175 (0.030)	-1.811 (<0.001)
<i>n</i> =	2987	3652

Individual level statistics

The results reported in the paper concern mainly group level and population level statistics. Here, we present the underlying individual level data.








Figure S6.3. Individual variation in reliance on historical information in helping behaviour. Panels show frequency distributions of fractions of decisions that were based on any information in the baseline (a), the noise (b) and the asymmetric (c) treatment. Rates were calculated over the entire sessions (on average 51 decisions). We observe considerable variation between individuals in their choice to access costly information.




Figure S6.4. Individual variation in use of direct experience and reputations. Panels show frequency distributions of relative importance of reputations in the baseline (a), the noise (b) and the asymmetric (c) treatment. Values were calculated as an individual's requests for indirect information, as a fraction of his total number of requests. We observe considerable variation in information use strategies. A large fraction of the population (about 50%) relies on either personal experience or reputations exclusively. In the presence of noise in indirect information, individuals tend to rely more on direct information (*i.e.* personal experience; b). When direct and indirect information are presented in unequal amounts (c), more people exclusively rely on reputations (c; Kruskal-Wallis multiple comparison $P < 0.05$). A GLMM with session replicate as random factor does not detect significant effect of individual reliance on reputations and overall helping rate ($p = 0.331$).




Figure S6.5. Total payoffs vary between individuals. In the experiment, individuals collected 5921 points on average (corresponding to €19.75, excluding the show-up fee). A linear mixed regression model with session as random factor reveals significant (negative) associations between helping rates and the total number of points collected. Information use (either direct or indirect, or both) has no significant effects.

Table S6.3. Correlations between individual characteristics ($n=60$ individuals per treatment). Requests for direct and indirect information are negatively correlated in each of the three treatments: when individuals request direct information, they are less likely to *also* request indirect information. Individuals that request more information tend to be more helpful. No relationships are observed between information use and payoffs.

Baseline treatment (a)

	direct				
indirect	-0.224 (0.085)	indirect			
both	-0.204 (0.118)	-0.016 (0.906)	both		
helping rate	0.334 (0.009)	0.253 (0.051)	0.112 (0.394)	helping rate	
payoff	-0.168 (0.201)	0.013 (0.921)	-0.048 (0.714)	-0.167 (0.202)	payoff
informed				0.488 (<0.001)	-0.146 (0.265)

Noise treatment (b)

	direct				
indirect	-0.322 (0.012)	indirect			
both	-0.028 (0.832)	-0.107 (0.415)	both		
helping rate	0.277 (0.032)	0.196 (0.134)	0.281 (0.03)	helping rate	
payoff	-0.026 (0.844)	0.048 (0.713)	-0.067 (0.613)	0.124 (0.346)	payoff
informed				0.515 (<0.001)	-0.018 (0.893)

Asymmetric treatment (c)

	direct				
indirect	-0.256 (0.048)	indirect			
both	-0.042 (0.749)	-0.245 (0.059)	both		
helping rate	0.092 (0.486)	0.228 (0.08)	0.056 (0.672)	helping rate	
payoff	0.044 (0.738)	0.083 (0.527)	-0.095 (0.471)	-0.075 (0.571)	payoff
informed				0.292 (0.024)	0.37 (0.778)

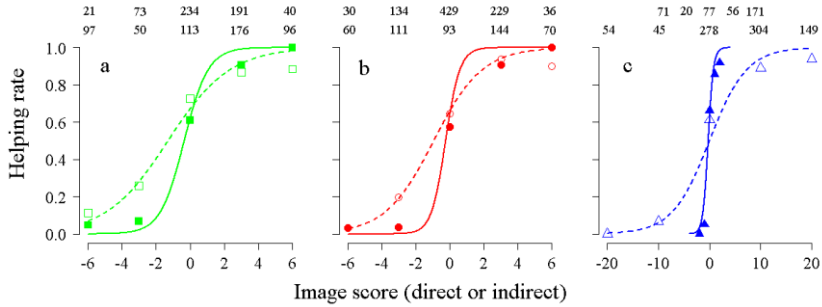


Figure S6.6. Fraction of helpful decisions preceded by requests for either direct (solid symbols) or indirect information (open symbols), in the baseline (a), the noise (b) and the asymmetric (c) treatment. Lines represent best least-square fits to the experimental data using logistic regressions (solid lines: behaviour towards the donor; dashed lines: behaviour towards others). The recipient's interaction history is calculated as the difference in the number of helpful decisions, *i.e.* as 'image score' (12). In (a) and (b), shown points are rounded towards the nearest modulus of 3. In (c) indirect information was rounded towards the nearest modulus of 10. The steeper slopes associated with direct information indicate that subjects react in a more extreme way to information coming from personal experience than to reputational information.

Robustness of the statistical findings

Table S6.4. To examine the robustness of the effects found in the regression model (Table 6.1 of the main text), we fitted another regression model where information on interaction histories was represented as the ‘image score’ (*cf.* Nowak and Sigmund 1998) rather than the fraction helpful decisions. Before entering the analysis, image scores for direct and indirect interactions were normalised as follows. In the baseline and noise treatment all image scores are divided by 6. In the asymmetric treatment direct and indirect image scores are divided by 2 and 20, respectively. Results of this generalised linear mixed model were in full agreement with the findings presented in the main text (including significant interaction between direct and indirect information). Another regression model fitted to the complete data set with all decisions – including those decisions not preceded by information requests – still detects significant effects of direct and indirect information, as well as the interaction between these types of information. Regression models based on alternative assumptions about the cumulative distribution function on the probability of helping (*i.e.* probit instead of a logit model) yield very similar patterns (not shown).

Glm regression results

	Estimate	<i>p</i>
Intercept	1.230	0.126
Period	5.276	0.049
Period ²	-7.391	0.002
Help donor	3.407	<0.001
Help others	2.067	<0.001
Help donor x help others	-1.838	0.001
Noise	-0.204	0.731
Asymmetric	-1.284	0.037

Instructions for the experiment

Welcome to this experiment on decision making. The experiment will last for approximately 90 minutes. During the session it is not allowed to talk or communicate with the other participants. If you have a question, please raise your hand and one of us will come to you to answer it. During this experiment you will make money. The amount you earn is dependent of your decisions and the decisions of others. At the end of the experiment the amount you have earned, plus a show up fee of 7 euros, will be paid to you in cash. These payments are anonymous; you will be paid individually in the reception room. Please remain seated at the end of the experiment until your desk number is called. We will not inform any of the other participants about your earnings. It is impossible for us to associate your desk number with your identity. You start out with an amount of 3000 points; 300 points are worth 1 euro.

Experimental procedure

1. Instructions
2. Quiz to verify if you understand the experiment
3. The experiment
4. Questionnaire
5. Payment

The experiment consists of at least 100 rounds. From round 100 upon, there is a chance of 90% that a new round starts. Every round, you will be paired with another participant in the room. Everybody stays anonymous; you will not be informed about the identity of the participant you are paired with.

The chance to be paired with a particular participant is for all participants the same in every round. Hence, the chance to be paired with the same participant twice in a row is very small.

Every round you will be assigned a role (A or B). You only have to make a decision when you are assigned role A. If you are assigned role B, you do not have to do anything.

Choice options

If you are assigned role A, you can choose between two alternatives. If you choose 'yellow', participant B, who is paired with you, receives 250 points. You lose 150 points. If you choose 'blue', the participant paired with you receives nothing, and you do not lose any points.

In summary:

Yellow has cost 150 for participant A. Participant B receives 250 points.

Blue yields 0 for both participants.

--- the following was specific to the three different treatments ---

a) Baseline treatment

Information

Before you make your decision, you have the opportunity to request information on the decisions of participant B in earlier rounds. By clicking the boxes you obtain a summary of the actions of participant B in up to **6 previous decisions in the role of A**. You can request two kinds of information.

1. Information about what participant B decided, in the role of A, when he was paired with **you**
2. Information about what participant B decided, in the role of A, when he was paired with **others**

This information is obtained by clicking the boxes. The information will be displayed on the screen like this:

1. Participant B decided in earlier rounds, when he was paired with **you** in role A: ... times yellow and ... times blue

2. Participant B decided in earlier rounds, when he was paired with **others** in role A: ... times yellow and ... times blue

Requesting information about participant B comes at a **cost**. This amounts to 5 points per requested information item. If you request both kinds of information, it costs 10 points. Only the participant clicking on the buttons obtains this information. The number of decisions of B you can see is 6 at maximum. If participant B has not been in the role of A yet, you will see a 0.

b) For noise treatment, this paragraph was added to the instructions of the baseline treatment

The information about what player B decided when he was paired with **you** is **perfectly reliable**. Information about what participant B decided when he was paired with **others** is **not perfectly reliable**. In one out of six cases, a 'blue' choice is displayed as 'yellow', or a 'yellow' choice is displayed as 'blue'. Thus, the information on what participant B did, when he was paired with others, is not completely reliable.

*c) For the asymmetric information treatment, the next paragraph replaced the 'Information' paragraph in sessions where indirect information was more abundant than direct information ****

Information

Before you make your decision, you have the opportunity to request information on the decisions of participant B in earlier rounds. By clicking the boxes you obtain a summary of the most recent actions of participant B. You can request two kinds of information.

3. Information about what participant B decided, in the role of A, when he was paired with **you** (maximally 2)

4. Information about what participant B decided, in the role of A, when he was paired with **others** (maximally 20)

This information is obtained by clicking the boxes. The information will be displayed on the screen like this:

3. Participant B decided in earlier rounds, when he was paired with **you** in role A: ... times yellow and ... times blue
4. Participant B decided in earlier rounds, when he was paired with **others** in role A: ... times yellow and ... times blue

Requesting information about participant B comes at a **cost**. This amounts to 5 points per requested information item. If you request both kinds of information, it costs 10 points. Only the participant clicking on the buttons obtains this information. If participant B has not been in the role of A yet, you will see a 0.

--- the final paragraphs were shown in all treatments ---

When participant A has made his decision, both participants will be informed about their payoffs in this round. This is the end of the round. In the next round you will be paired with another randomly drawn participant. The roles of A and B are randomly assigned.

Please remain seated at the end of the experiment, until we call your desk number. You will be paid in the reception room individually.

--- end of instructions ---

*** We ran two pilot sessions for the asymmetric treatment with instructions stating that with a request for direct or indirect information, a donor would obtain the recipient's decisions in the past 44 rounds of the experiment. In these two sessions cooperation levels were slightly higher, perhaps because subjects were more primed to be the fact that cleaning their record would take a long time. Patterns of requests for direct and indirect information were very similar to the results reported in the main text.

Chapter seven

Diversity of morals in human cooperation

Lucas Molleman
Violet Swakman
Aljaž Ule
Martijn Egas

Manuscript in preparation

Abstract

Human cooperation can be sustained when help is channelled towards those with a good reputation. Evolutionary models of ‘indirect reciprocity’ show that conditioning help on recipients’ recent behaviour alone cannot stably support cooperation. Stable reputation-based cooperation requires that individuals consider the motivation behind the past behaviour of their interaction partners, and distinguish justified from unjustified decisions to help. Empirical evidence that people actually use such strategies is weak. Here we show experimentally that people do consider their peers’ motivations, and reward those individuals who refuse to help defectors. Moreover, we report pronounced differences in individual strategies: a majority bases their decisions to help merely on the past actions of their interaction partners, but a substantial proportion of individuals additionally considers the motivation behind these actions.

Introduction

Cooperation is a puzzle to scientists in many domains of science. Why would an individual provide costly help to a potential competitor? For cooperation to thrive, specific mechanisms are needed to prevent defectors reaping the benefits from cooperation without paying the costs. Such mechanisms either enforce cooperation through punishment (Yamagishi 1986; Ostrom et al. 1992), or channel help towards cooperators through a form of assortment (*e.g.* kin selection (Hamilton 1964), group selection (Maynard Smith 1964) or partner choice (Noë and Hammerstein 1994)). When individuals interact repeatedly, scratching someone's back can pay off when he can scratch yours at a later point in time (Trivers 1971). Such 'direct reciprocity' is widespread among humans and is believed to be present in some animal species as well (Dugatkin and Reeve 1997; Clutton-Brock 2009).

Human cooperation can also be supported by *indirect* reciprocity (Darwin 1871; Alexander 1987). Frequently, interactions between individuals are observed by others, and the outcome is often communicated through gossip. Over time, people can build up a good reputation when they tend to be nice and helpful towards others, or a bad reputation when they tend to be nasty and self-centred. It can pay off to provide costly help to someone, when a good reputation is more likely to attract help from others.

The dynamics of reputation-based strategies has been modelled using 'image scores' that reflect the helpfulness of an individual in recent interactions (Nowak and Sigmund 1998). Theoretical analyses suggest that in a randomly interacting population, cooperation can be promoted by strategies that condition their decisions to help on the image scores of their recipients (Nowak and Sigmund 1998; Nowak and Sigmund 2005). However, a population of individuals using such strategies can be invaded by mutants that do not reciprocate but merely uphold their own reputation in order to keep receiving help. In turn, the resulting loss of reciprocation paves the way for defectors to take over (Leimar and Hammerstein 2001). More sophisticated strategies examine the previous actions of their interaction partners more deeply, in which one's good reputation is only damaged when one fails to help a cooperator, whereas

refusing to help defectors does not affect one's good 'standing' (Sugden 1986). These strategies ensure that individuals attain a good reputation by channelling help towards cooperators and away from defectors, thereby stabilizing cooperation (Panchanathan and Boyd 2003; Brandt and Sigmund 2004; Ohtsuki and Iwasa 2004; Ohtsuki and Iwasa 2006).

Employing such sophisticated strategies requires, on top of the (first-order) information about a recipient *R*'s helpfulness, also (second-order) information about *R*'s recipients at the time he made his decisions to help. Experiments on indirect reciprocity indicate that in human groups, the availability of first-order information can lead to considerable levels cooperation (Wedekind and Milinski 2000; Wedekind and Braithwaite 2002; Seinen and Schram 2006), but that the presence of second-order information increases levels of cooperation even more (Bolton et al. 2005). It has been argued that in real-life interactions, people unlikely take into account both first- and second-order information since it is cognitively too taxing (Panchanathan and Boyd 2003). This view is supported by experimental studies on indirect reciprocity failing to find clear indications for strategies that follow second-order information in decision making (Milinski et al. 2001; Ule et al. 2009).

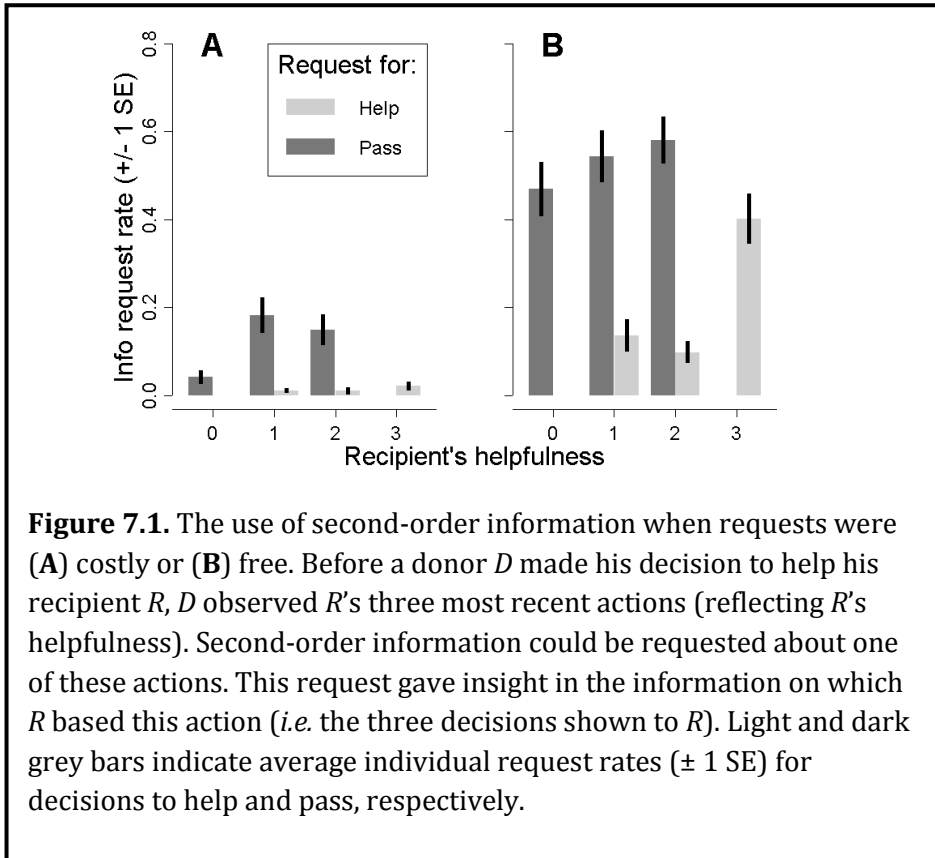
However, these experiments were not primarily designed to directly assess the use of second-order information, and the rich sets of historical information available to participants likely disfavoured the use of more sophisticated strategies. This leaves us with a situation in which on the one hand theory predicts that strategies accounting for second-order information can more stably support cooperation through indirect reciprocity, whereas on the other hand no experimental evidence decides whether such information plays a role in human interactions.

To shed light on this issue, we conducted an experiment assessing the use of first- and second-order information in indirect reciprocity. A total of 140 subjects were invited in the computer laboratory, and in groups of ten these subjects played 100 rounds of an 'indirect helping game' (Seinen and Schram 2006). In each round, pairs were randomly formed, and one of the two (the 'donor') could decide to either *help* the other (the 'recipient'), or to *pass*. Helping incurred a cost of 200 points to the donor, increasing the recipient's payoff with 250 points (where 300 points = 1 euro);

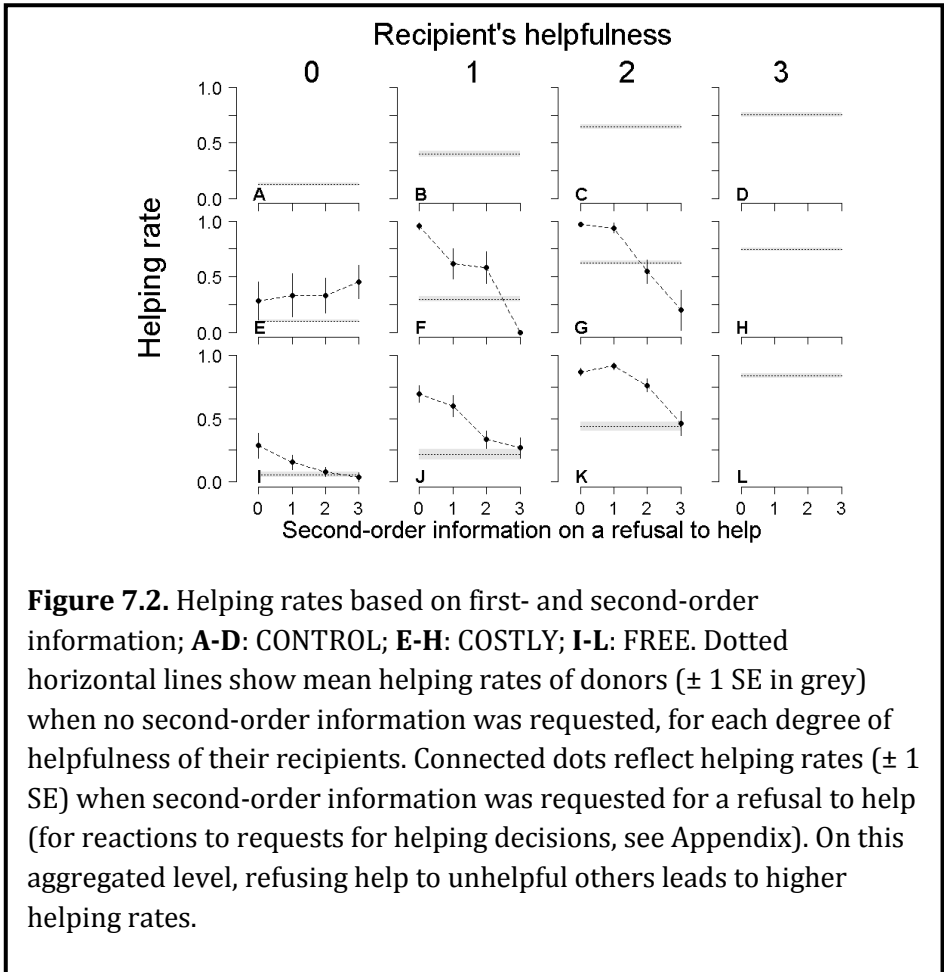
passing did not lead to any change in payoffs. Before making their decision to help, donors always received information about the three most recent actions of their current recipient. This (first-order) information was given in random order. Donors could choose to request second-order information about one of these three decisions by clicking on a button; this (second-order) information consisted of the information the recipient had when he or she made this decision. We imposed three experimental conditions: second-order information was either *i*) unavailable (CONTROL; four replicate groups); *ii*) available at a cost of five points (COSTLY; six replicate groups); or *iii*) available for free (FREE; four replicate groups, making a total of 14 independent observations). For details of the experimental design, including paper instructions and screenshots, see section 3 of the Appendix.

Results

The overall helping rate was 55%, which is comparable to helping rates in other experiments on indirect reciprocity (Wedekind and Milinski 2000; Bolton et al. 2005; Seinen and Schram 2006; Ule et al. 2009), and did not differ significantly between treatments. The dynamics of cooperation levels over time and detailed statistics are presented in sections 1.1 and 1.2 of the Appendix. Subjects requested second-order information more frequently when this information was free, and when first-order information showed intermediate levels of helpfulness (Figure 7.1; logistic generalized linear mixed model: $P < 0.001$ for both effects; see Appendix, section 1.3, for details). Overall, requests were strongly biased towards recipients' refusals to help others (*i.e.*, decisions to pass; χ^2 test: $P < 0.001$; compare the dark and light grey bars in Figure 7.1). Below we focus our analysis of the use of second-order information on these decisions to pass only. In section 1.4 of the Appendix, we analyse the use of second-order information on recipients' decisions to help.

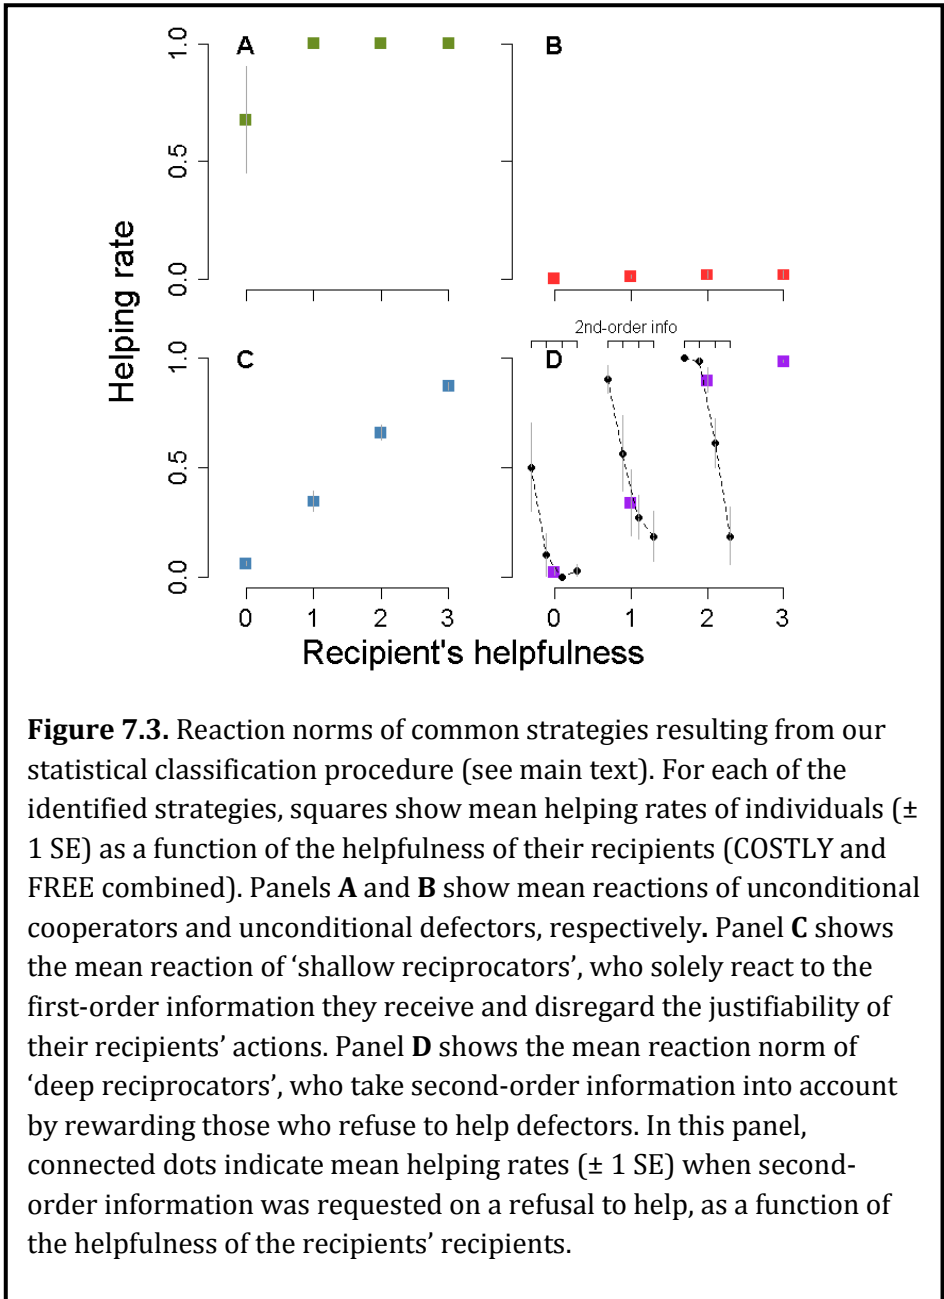


On average, subjects reacted to first-order information by rewarding helpful individuals and refusing to help defectors (Figure 7.2, dotted horizontal lines). When requested, second-order information had a strong effect on helping behaviour: helping rates are substantially higher when a decision to pass was targeted at an unhelpful individual (Figure 7.2E-G and I-K, connected dots tend to decrease; logistic generalized linear mixed model: $P < 0.001$, see Appendix, section S7.1.2, for details). This result indicates that people care whether or not the defections of their interaction partners in the past were justified.



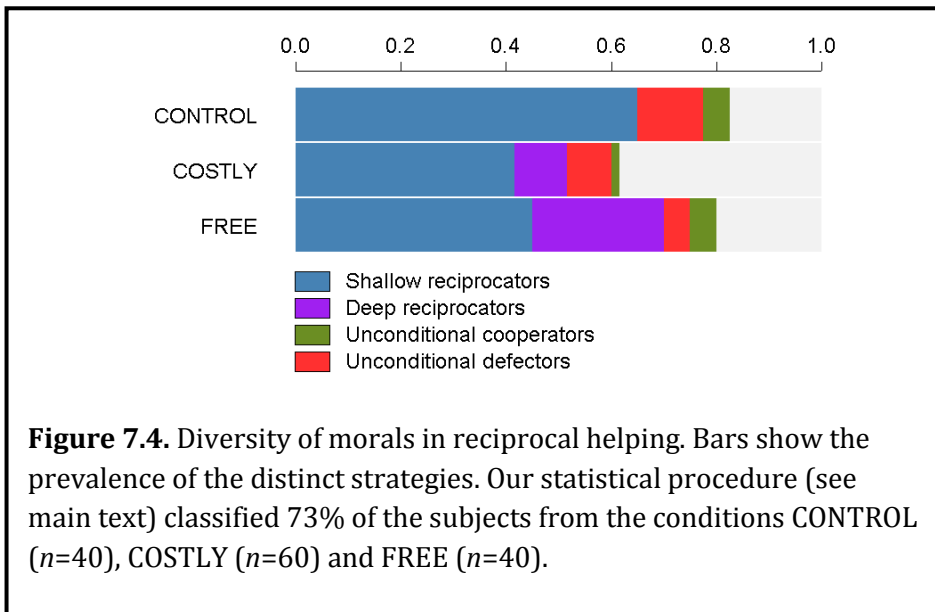
Recent experimental evidence suggests that even within demographically homogeneous populations, humans differ amongst each other in terms of strategic decision making (Kurzban and Houser 2005; Engelmann and Fischbacher 2009, Chapters 4 and 5 of this thesis). Moreover, such diversity can have a strong impact on group dynamics and the outcome of social interactions (Fischbacher and Gächter 2008; Wolf et al. 2008; McNamara and Leimar 2010, Chapter 5 of this thesis). To gain insight in the individual variation underlying the results presented in Figure 7.2, we analyzed individual reactions to first- and second-order information by fitting statistical models to the decisions of each of our 140 subjects separately. We labelled as ‘reciprocators’ those subjects for which positive

first-order information on recipients' helpfulness had a significantly positive effect on helping rates. When a reciprocator requested second-order information on at least five occasions, we fitted a separate model to these decisions to test for sensitivity to motivations behind refusals to help (see Appendix, section S7.2.2, for details). When subjects' helping rates significantly decreased as their recipients failed to help more cooperative individuals, they were labelled as a 'deep reciprocator'. When no such effect was detected or second-order information was requested fewer than five times, a subject was labelled as a 'shallow reciprocator'. Subjects not labelled as reciprocators were categorised as unconditional cooperators or defectors, in case they helped in more than 90% or less than 10% of the cases, respectively. This procedure classified 73% of the participants, leaving 27% unclassified. A similar statistical procedure revealed that about 5% of the unclassified subjects strategically maintained their reputations, basing their decisions on their own previous actions exclusively (see Appendix, section S7.2.3, for details). Figure 7.3 shows of the classified types the average response to first- and second-order information.



The prevalence of each of these distinct types is shown in Figure 7.4. Deep reciprocators, who evaluate the justifiability refusals to help others,

present a substantial fraction of the total population when second-order information was available. Whereas costs of second-order information decreased overall information request rates (see Figure 7.1), it does not lead to significant differences in the prevalence of deep reciprocators (Fisher-exact test: $P=0.11$). Shallow reciprocators, who more readily help cooperators based on first-order information only, are the most common types. Other types are relatively rare. Identified strategies do not strongly differ in their performance. The only significant differences in performance was found between the two unconditional strategies: defectors achieve higher payoffs than cooperators (Tukey's test on linear mixed model: $P=0.028$, see Appendix, section 2.1, for details).



Discussion

In summary, our experiment provides evidence that cooperative decisions are affected by the motivations of interaction partners: refusals to help defectors are likely to be rewarded. However, this (potentially costly)

virtuous behaviour is unevenly distributed in the population, as highlighted by the mixture of shallow and deep reciprocators we observe. These results support the emerging view that cooperative strategies do not only strongly differ between cultures (Henrich et al. 2001; Herrmann et al. 2008; Gächter et al. 2010), but also within populations of demographically similar individuals (Kurzban and Houser 2005; Engelmann and Fischbacher 2009; Chapters 4 and 5 of this thesis).

This variation in reciprocal strategies can have marked consequences on the emergence and stability of cooperation in a group. For instance, a small fraction of deep reciprocators might suppress free riding, thereby sustaining cooperation. More stable cooperation can be established when individuals can actively choose with whom to interact (Rand et al. 2011). If individuals could estimate the reciprocal strategies of others and assort accordingly, clusters of deep reciprocators may stably maintain reputation-based cooperation, whereas cooperation among shallow reciprocators may break down in the face of encroaching defectors.

Evolutionary models can help understand the emergence of strategic polymorphisms of the kind we report here (Wolf et al. 2007; Wolf and Weissing 2012). In particular, the benefits of social sensitivity can depend on the variation in cooperative tendencies in a population (Wolf et al. 2008; McNamara and Leimar 2010; Wolf et al. 2011): whereas deep reciprocators can sustain stable reputation-based cooperation in the presence of defectors, their costly social sensitivity is disfavoured by selection once cooperation has established in the population. Cooperative – but less sensitive – shallow reciprocators can then increase in frequency, making the population more vulnerable for invasion of defectors.

Although the relatively rich information conditions of our experiment allow for complex and perhaps even realistic strategies, we employed a strong simplification of the concept of a ‘reputation’. In reality, reputations build up over various contexts of interaction, and factors like information reliability play a role. The mechanisms of information transfer are crucial: one may directly observe interactions between individuals and change one’s esteem of them, but reputations also propagate indirectly through

gossip (Sommerfeld et al. 2007; Sommerfeld et al. 2008). Gossip is prone to error, lies, and strategic manipulation, and therefore second-order information could be less reliable than first-order information (Ohtsuki and Iwasa 2004; Panchanathan 2011). Such unreliable or manipulated information would impair the efficiency of deep reciprocity in supporting cooperation.

Individuals interacting in small groups often base their decisions not only on reputations, but also on previous encounters with others; in this case cooperative behaviour may depend on the interplay between direct and indirect reciprocity (Roberts 2008; Molleman et al. 2013). From this perspective, our understanding of human cooperation through reciprocal helping could benefit from theoretical and empirical investigations of how reputational information is weighted with reliability and integrated with information from other sources such as direct experience.

Appendix

This Appendix comprises 3 sections. In **Section 1**, we show mean cooperation rates in each of the replicate sessions of our experiment, and outline the statistical comparison of cooperation levels between the three conditions (CONTROL, COSTLY, and FREE) we imposed (S7.1.1). Also, we present results of a set of statistical models assessing which factors influence information requests (S7.1.2), and how first- and second-order information affect decision making on the aggregate level (S7.1.3). We present analyses of reactions to second-order information about refusals to help (on which the main text is focused), but also show average reactions after requests for helping decisions. **Section 2** starts with a comparison of the payoffs achieved by each of the four strategy types distinguished in the main text (S7.2.1). Further, we offer a detailed description of the statistical models fitted to decisions of each individual (S7.2.2). These models underlie the classification procedure which led to the strategy types presented in Figure 7.3 and 7.4 of the main text. Subsequently, we present a more detailed analysis of individual decisions by fitting a different statistical model to the data (S7.2.3). Finally, **Section 3** presents details of our experimental design and procedures (S7.3.1 and S7.3.2), including (translations of) the paper instructions given to the subjects, as well as screenshots of the decision making environment (S7.3.3).

S7.1. Cooperation levels in the experiment and aggregate responses to second-order information

This section starts by describing the dynamics of cooperation in our experimental sessions. We compare helping rates between treatments by fitting statistical models to the data. These models also assess the effects of reactions to first- second order information. We further assess which factors influence decisions to request second-order information. Finally, we show the aggregate reactions to second-order information on decisions to help.

S7.1.1. Dynamics of cooperation




Figure S7.1.1. Dynamics of helping over time. Panels show mean helping rates for each of the three experimental conditions for each of the replicate sessions. Symbols show 10-round averages of helping rates over the course of time. Each line represents one replicate session; $n=50$ per data point (10 rounds x 5 decisions per round).

S7.1.2. Reactions to first- and second-order information: aggregate-level statistics

We fitted three different models to the data to assess the effects of our experimental conditions (CONTROL, COSTLY, and FREE) and the contents of first- and second-order information displayed to our subjects. Each of these models is a logistic generalized linear mixed model fitted to the decisions to help, and each takes as fixed factors 'Period' (included because cooperation rates tend to decrease over time, see Figure S7.1.1), and 'Experimental condition'. Subject nested in replicate group was used as random factor. The effect estimates of these three models are presented in Table S7.1.2.1. Model 1 tests whether cooperation levels differed between the conditions CONTROL, COSTLY and FREE, and includes all ($n=7000$) decisions. This model reveals that the availability of second-order information slightly increased cooperation levels, but this effect was not significant ($P = 0.571$ and 0.091 for COSTLY and FREE, respectively). Model 2 is the same as Model 1, but we added one element: the contents of first-order information (*i.e.* the number of helping decisions out of the recipient's three most recent actions), to assess how such information affects decisions to help. The significantly positive estimate for first-order information suggests reciprocal motives of our subjects: help was more likely as first-order information showed more helping decisions of the recipient. Model 3 considers only those ($n=924$) decisions in which donors had requested second-order information about a recipient's refusal to help (*i.e.*, a *pass* decision). This model assesses the effects of second-order information, while controlling for the effect of first-order information. The significantly negative estimate for second-order information indicates that help was more frequent when a decision to pass was directed at an unhelpful individual. This suggests that people tend to take into consideration the motivations behind the actions of their interaction partners.

	Model		
	1	2	3
Intercept	0.772 (0.028)	0.475 (0.174)	1.089 (<0.001)
Period / 100	-2.002 (<0.001)	-1.700 (<0.001)	-0.006 (0.058)
COSTLY	0.25 (0.571)	0.25 (0.570)	
FREE	0.830 (0.091)	0.808 (0.095)	-0.473 (0.039)
First-order info		0.056 (<0.001)	0.312 (<0.001)
Second-order info			-0.189 (<0.001)

Table S7.1.2.1. Effect sizes of three models fit to decisions to help.
 Estimated effects are shown for Model 1, 2 and 3, respectively, with *P* values in parentheses (see above text for model descriptions).

S7.1.3. Information requests

Figure 7.1 of the main text shows the distribution of information requests in our experiment. To assess the effects of information cost and the content of first-order information (*i.e.*, the helpfulness of recipients), we first coded the helpfulness of recipients as ‘intermediate’ (helped 1 or 2 times out of the most recent three actions), or ‘extreme’ (helped 0 or 3 times out of the most recent three actions). Subsequently, we fitted a logistic generalized linear mixed model to individual decisions to request information, including ‘information cost’ (COSTLY or FREE) and ‘helpfulness of recipient’ (intermediate or extreme) as fixed factors, and ‘subject’ as random factor. This model detects a strong negative effect of information cost ($P < 0.001$), indicating that information requests are significantly more frequent when they were for free. We also observe a strong effect of the extremity of recipients’ helpfulness on information requests ($P < 0.001$): requests are significantly more frequent when first-order information shows intermediate levels of helpfulness. A plausible interpretation of this finding is that people use second-order information at times when first-order information is inconclusive.

S7.1.4. Reactions to second-order information on decisions to help

The main text focuses on reactions to second-order information on recipients' refusals to help (*i.e.*, a decision to *pass*). This is motivated by the fact that requests were strongly biased towards these decisions to pass (87.5% and 63.5% of the requests, in COSTLY and FREE, respectively). Here, we focus on cases where subjects requested information on a decision to *help*, and analyze reactions to second-order information along the lines of the analysis whose results are presented in Figure 7.2 of the main text.

Figure S7.1.4.1. shows that the effects of second-order information are much less pronounced when they are about decisions to help (as opposed to decisions to pass). This has two reasons. First, in the COSTLY condition, the analysis is based on only 36 cases in which second-order information on a helping decisions was requested. In the FREE condition we have 386 of such cases, but these largely correspond to situations in which a recipient had chosen to help in all three of his most recent decisions (see Figure 7.1B of the main text, bars to the right hand side). In these cases, helping rates were generally very high. Since second-order information was available without any monetary cost, subjects might have requested this information not to account for it in their current decision to help, but rather in order to explore the general cooperativeness in their social environment. Such behaviour would dilute the effects of second-order information in this analysis. However, we do find some indications that people are considering the motivations of their interaction partners, also after a request for information on a decision to help. From panels H and K of Figure S7.1.4.1, we observe that helping decisions towards cooperators are rewarded more (the points tend to go up with more positive content of second-order information; *i.e.* the helpfulness of the recipients' recipients).




Figure S7.1.4.1. Helping rates based on first- and second-order information; **A-D**: CONTROL; **E-H**: COSTLY; **I-L**: FREE. As in Figure 2 of the main text, dotted horizontal lines show mean helping rates of donors (± 1 SE in grey) when no second-order information was requested, for each degree of helpfulness of their recipients. Connected dots reflect helping rates (± 1 SE) when second-order information was requested for a helping decision. Effects of second-order information are less pronounced, although panels **H** and **K** suggest that more help is given when helping decisions are 'justified'.

S7.2. Payoffs of strategies and statistical models fit to decisions of individuals

Figure 7.3 and 7.4 of the main text show that the subjects in our experiment used distinct strategies in making their decisions. In this section, we analyse the performance of these strategies in our experiment. Further, we give a detailed description of the statistical procedure used to classify the subjects. We also describe an alternative classification based on a more complicated model.

S7.2.1. Performance of identified strategies

To assess the performance of the four strategy types identified in the main text, we calculated for each subject their average payoffs as a donor and as a recipient separately. By subtracting the costs in the role of donor role from the benefits in the role recipient, we obtain weighted average payoffs per round for each individual. Figure S7.2.1.1 shows for each of the identified strategy types these mean payoffs (± 1 SE). We do not observe strong payoff differences between the different strategies in our experiment: significant differences (at the 5% level) in payoffs only occur between unconditional defectors and cooperators (Tukey's test on a linear mixed model fit to payoffs of different types with 'session' nested in 'treatment' as random factor: $P=0.028$).




Figure S7.2.1.1. Payoffs of strategy types in our experiment. For each individual we calculated its weighted average payoff per round by subtracting its mean costs as a donor from its mean benefits as recipient. Bars show for each type the population average payoffs for each strategy type (± 1 SE).

S7.2.2. Statistical classification procedure

We analysed the use of first- and second-order information of each of our 140 subjects separately. This analysis underlies the classification of types presented in the main text. First, we tested whether the cooperation rate of an individual was either higher than 0.9, or lower than 0.1. In these cases, a subject was classified as an unconditional cooperator or defector, respectively. When a subject's cooperation rate was between 0.9 and 0.1, we tested whether the content of first-order information (*i.e.*, the recipient's helpfulness in recent interactions) had a significant effect on helping behaviour. To this end, we fitted a logistic generalized linear model to the decisions to help, while including – apart from the content of first-order information – ‘Period’ to control for time effects on cooperation rates. If this model detected a significantly positive effect of first-order information (with $P < 0.05$), a subject was classified as a reciprocator. When reciprocators requested second-order information about their

recipients' refusals to help at least five times, we tested whether this information affected their decisions to help by fitting a Bayesian generalized model to these decisions, using the content of second-order information (*i.e.*, the helpfulness of the recipients' recipients) as the only predictor. We used Bayesian regression to avoid problems of linear separation in our data. When this model detected a significantly negative effect of second-order information (again with $P < 0.05$) – so that refusing to help defectors was rewarded more, but refusing to help cooperators was rewarded less – a subject was classified as a 'deep reciprocator'. When no such effect was found (or when second-order information was requested fewer than five times), a subject was classified as a 'shallow reciprocator'. Fisher-exact tests comparing frequencies of help in case of negative (0 or 1 times help) or positive (2 or 3 times help) second-order information led to the exact same results in distinguishing shallow from deep reciprocators. Also, in a questionnaire following the experiment, the majority of individuals classified as 'deep reciprocator' indicated that they followed second-order information in their decisions to help. Our statistical procedure classified 102 out of our 140 (73%) of our participants.

S7.2.3. Alternative classification

In games of indirect reciprocity, it can pay off to maintain a certain level of cooperativeness to attract help from others. Such a strategy does not require paying attention to the helpfulness of your interaction partners in making decisions to help. Such non-reciprocal types can destabilize cooperation (Leimar and Hammerstein 2001; Panchanathan and Boyd 2003). To test whether individuals in our experiment used such a strategy, we ran the same analysis as above, but we added one element. We added to the statistical model fit to the decisions of each individual his own most recent two decisions. 'Hypocritical' individuals that strategically maintain a level of cooperation without reciprocating, do not react to the first-order information about recipients' helpfulness, but only help when they refrained from helping in recent interactions. Individuals with a more complex strategy, reacting (positively) to both their recipients' helpfulness and (negatively) to their own recent helpfulness, were

assigned to a separate category. With this more advanced procedure, we classified 78% of our participants.

In analogy to the main text, Figure S7.2.3.1. presents the average reactions of each of the classified types, and Figure S7.2.3.2. shows their relative frequencies in each of the three experimental conditions. First of all, we observe that the results are very similar to those presented in Figure 7.3 and 7.4 of the main text. In addition, we can see that some individuals tend to increase their helpfulness after having recently refrained from helping. Types that account for their own image ('hypocrites' and 'complex' types) are most frequent in the CONTROL treatment, in which no second-order information was available. This made it impossible for others to assess the motivations behind the decisions of interaction partners. For the 'complex' types, despite the fact that the statistical model detected significant negative effects of own helpfulness, the recipients' helpfulness seems more important in decision making. We do not observe significant differences (at the 5% level) between strategies in terms of their total payoffs, suggesting that these strategies perform equally well in our experimental setting (Tukey's test on a linear mixed model with 'session' nested in 'treatment' as random factors).

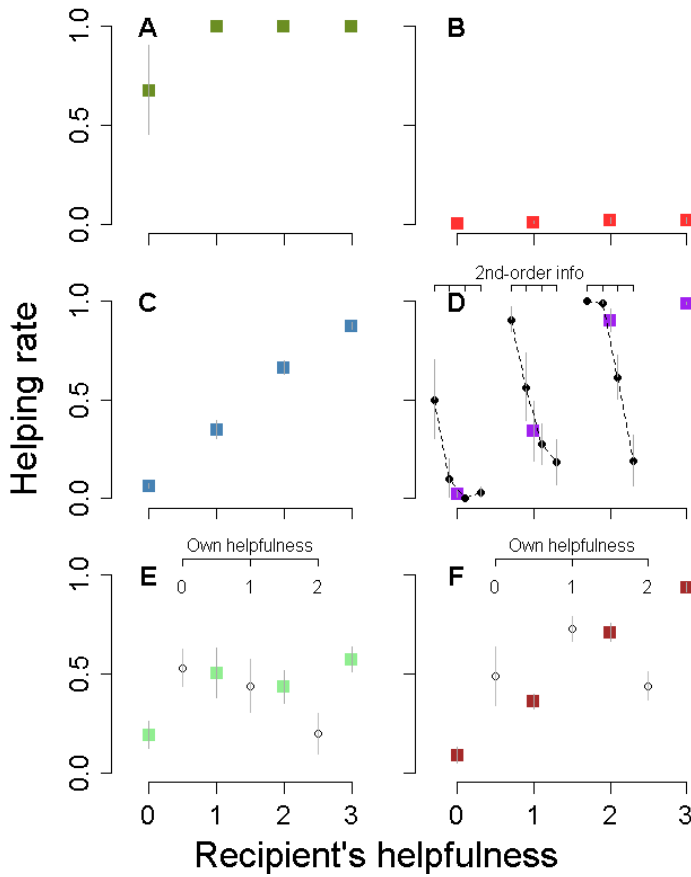


Figure S7.2.3.1. Reaction norms of common strategies resulting from the alternative classification procedure. For each of the identified distinct strategies, squares show mean helping rates of individuals (± 1 SE) as a function of the helpfulness of their recipients (COSTLY and FREE combined). Panels A-D as in Figure 7.3 of the main text. Open circles in panels E and F indicate mean helping rates of individuals (± 1 SE) as a function of their own recent helpfulness. Panel E shows the reaction norm of ‘hypocrite’ individuals that based their decisions on their own recent helpfulness, maintaining an appearance of being helpful, without reciprocating. Panel F shows the reaction for individuals with more complex strategy, basing decisions on both the helpfulness of both themselves as well as that of their recipients.

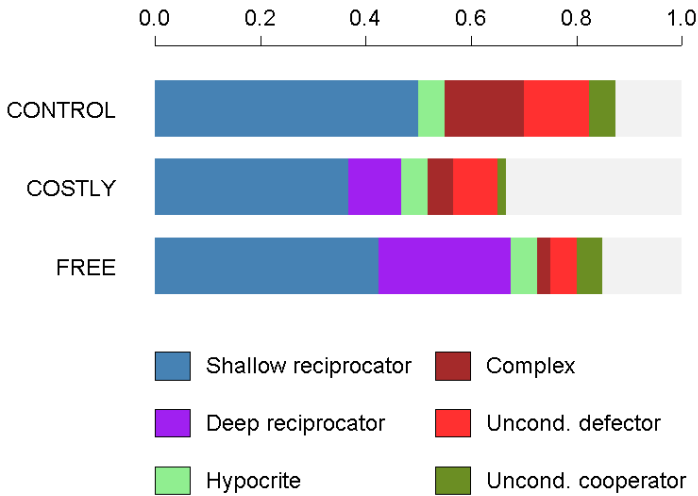


Figure S7.2.3.2. Frequency of strategies from the alternative classification procedure. Shallow reciprocators condition their decisions to help on the helpfulness of their recipients. Deep reciprocators additionally consider the motivations behind their recipients’ defections. Hypocrites tend to help only when their recent history would paint a negative picture for potential donors. Such maintenance of a cooperative image also play a role for complex individuals, but they also consider their recipients’ helpfulness. Unconditional cooperators and defectors are unaffected by any such information.

S7.3. Details on the experimental design

This section describes general experimental procedures and gives more details on the experimental design. Also, we here present the paper instructions given to the subjects and show screenshots of the decision making environment.

S7.3.1. Experimental procedures

We ran 14 replicate sessions at the CREED laboratory at the University of Amsterdam and the Sociology laboratory of the University of Groningen. In total, $n=140$ subjects attended the sessions, participating in groups of 10 (CONTROL: four replicates; COSTLY: six replicates; FREE: four replicates). Subjects were mostly undergraduate students from the social sciences, economics and biology. Experimental sessions lasted around 90 minutes, in which subjects earned €14.60 (\$19.25) on average according to their performance, excluding a show-up fee of €7 (making a total of €21.60 = \$28.50). The experiment was conducted using Z-Tree (Fischbacher 2007). Code is available upon request.

In the reception room, participants received a randomly chosen number corresponding to a desk in the computer laboratory. As soon as the subjects entered the laboratory, they were no longer allowed to communicate with each other. Paper instructions (see below) were distributed and were read out loud by one of the experimenters. Before the experiment itself started, subjects had to fill out a brief quiz to check their understanding of the experiment. Sessions finished with a questionnaire including items about personal demographic background. Subjects were paid individually in the reception room.

S7.3.2. Indirect helping game

In groups of ten, subjects interacted in 100 rounds of an ‘indirect helping game’. Each round, pairs were randomly formed and within each pair, the

roles of donor and recipient were assigned (in the experiment referred to as role A and B). Interactions were anonymous. Donors decided between two options, either clicking on 'blue' (*i.e.* help; thereby increasing the score of his recipient with 250 points and decrease their own score with 200 points, where 300 points correspond to 1 euro) or 'purple' (*i.e.* pass; not changing the payoffs of either). Recipients did not have to do anything, and a waiting screen was displayed requesting to wait for the decision of the coupled donor.

Before making their decisions, the three most recent decisions of the recipient were always shown to the donors (see screenshot of Figure S7.3.2.1). These decisions were displayed in random order. In the treatments COSTLY and FREE, donors had the option to request second-order information about one of these three decisions, by clicking the button 'more information' below one of the decisions (Figure S7.3.2.1). This second-order information contained the three most recent decisions of the player this decision was directed at (Figure S7.3.2.2). In treatment COSTLY the second-order information was available at a small cost (5 points; €0.017) in the treatment FREE this information was free. After deciding whether or not to request second-order information, donors made their decision to help or pass. In the first rounds, when not all history information was complete, the missing data was displayed as a hyphen.

At the end of each round, the donor's decision and the payoff consequences were shown to both the donor and his recipient.

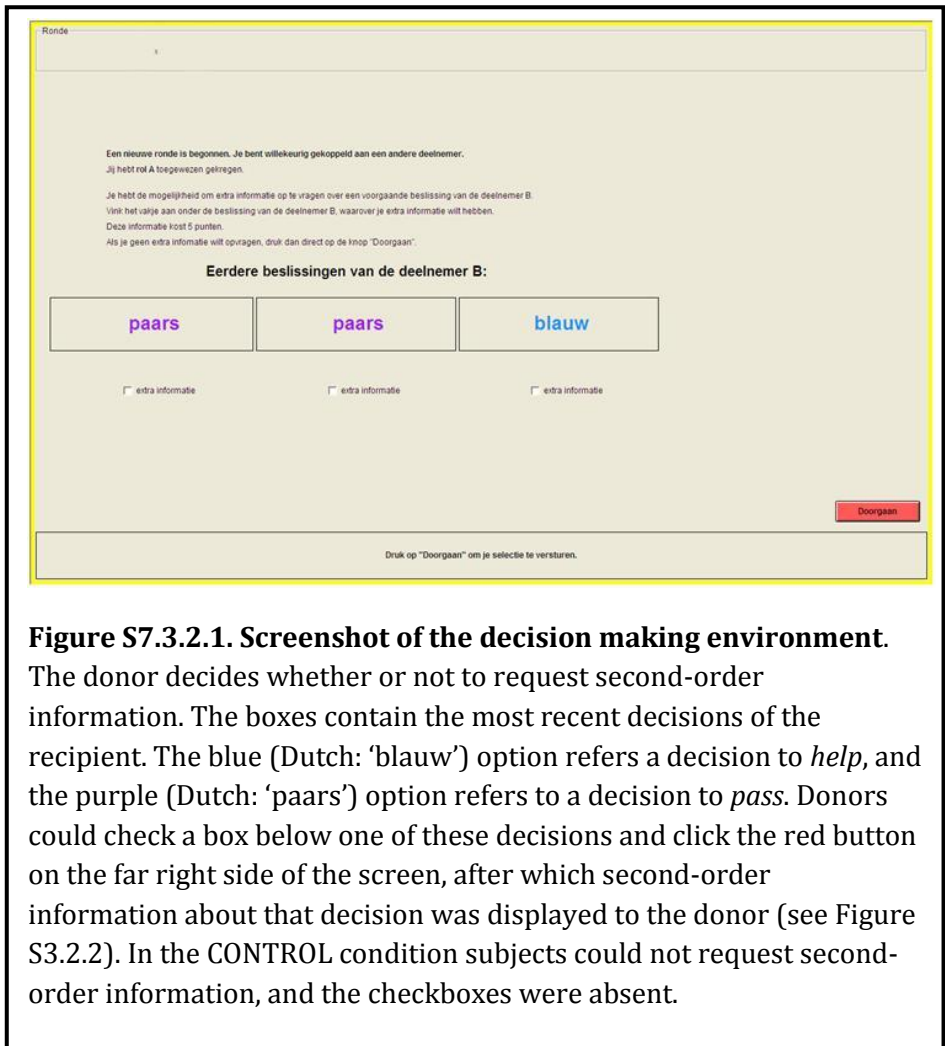


Figure S7.3.2.1. Screenshot of the decision making environment.

The donor decides whether or not to request second-order information. The boxes contain the most recent decisions of the recipient. The blue (Dutch: 'blauw') option refers a decision to *help*, and the purple (Dutch: 'paars') option refers to a decision to *pass*. Donors could check a box below one of these decisions and click the red button on the far right side of the screen, after which second-order information about that decision was displayed to the donor (see Figure S3.2.2). In the CONTROL condition subjects could not request second-order information, and the checkboxes were absent.

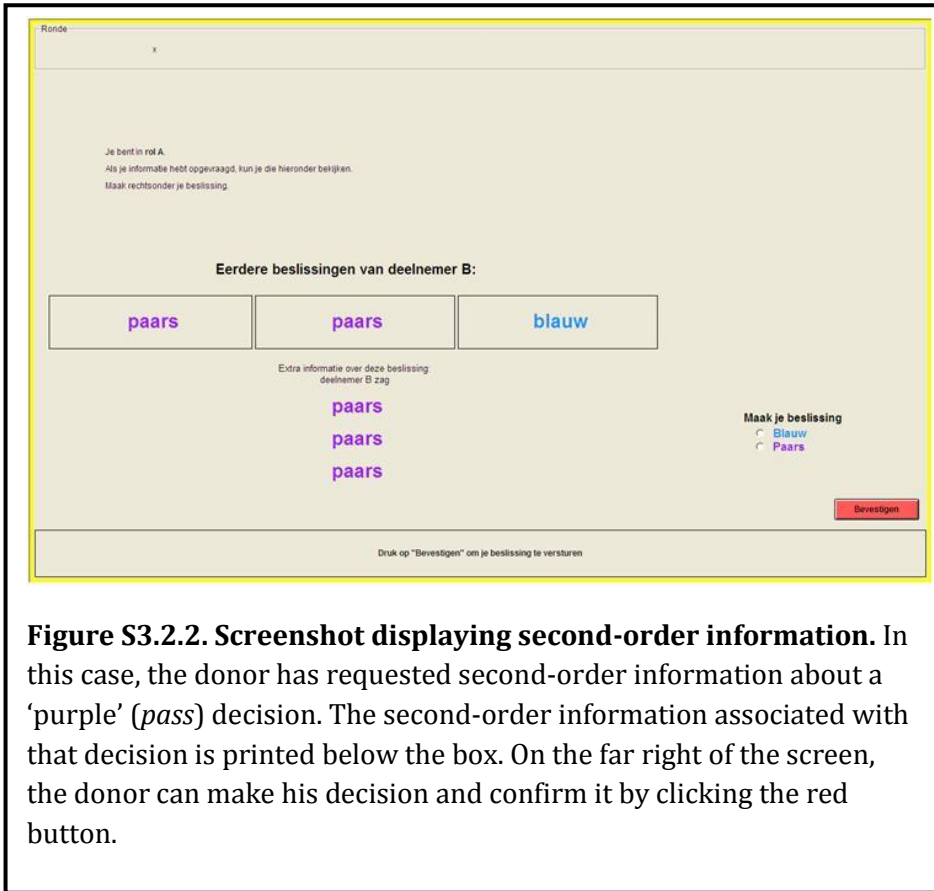


Figure S3.2.2. Screenshot displaying second-order information. In this case, the donor has requested second-order information about a ‘purple’ (*paars*) decision. The second-order information associated with that decision is printed below the box. On the far right of the screen, the donor can make his decision and confirm it by clicking the red button.

S3.3. Instructions on paper

Subjects received instructions on paper. For each of the three experimental conditions, instructions were largely identical, but differed with respect to the availability (and costs) of second-order information. In the below printed text, we will indicate which parts are identical for each experimental condition and which parts are specific to either of them. Original instructions were in Dutch, the below text is a translation.

=== *Instructions* ===

Introduction

Welcome to this experiment on decision making. The experiment lasts for about 90 minutes. During the session it is not allowed to talk or communicate with the other participants. If you have a question, you can raise your hand and one of us will come to you to answer it.

During the experiment you will make money. Your earnings depend on your decisions and the decisions of others. At the end of this session, the money you earned, plus a 7 euro show-up fee, will be paid to you in cash. These payments are anonymous, you will be paid individually in the reception room. After the experiment has ended, please remain seated until we call your desk number. None of the other participants will be informed about your earnings. We cannot connect your desk number to your identity. You start out with 3000 points, and 300 points are worth 1 euro.

Setup of this session

1. Instructions
2. Quiz to check if you understand the experiment
3. The experiment itself
4. Questionnaire
5. Payment

The experiment consists of 100 rounds and takes place in groups of 10. In each round, you will be paired with one of the nine other participants in your group. Everyone remains anonymous; you will not receive information about the identity of the participant you are paired with.

In each round again, a role (A or B) is randomly assigned to you. You only have to make a decision when you are assigned role A. When you are assigned role B, you do not have to do anything.

Making decisions

When you are assigned role A, you can choose between two alternatives. When you choose 'blue', player B that is paired with you received 250

points. You lose 200 points. When you choose 'purple', the player that is paired to you receives nothing, and you do not lose any points.

In summary:

Blue costs participant A 200 points. Participant B receives 250 points.

Purple yields 0 points for both.

=== *The following part was specific to the condition where second-order information was not available (CONTROL) ===*

Information

When you are taking your decision, the screen will show the three most recent decisions that the participant in role B has taken, when he was in the role of A.

These three decisions of the participant in role B are displayed in three boxes on your screen. These boxes are given in random order. When B has not yet been in the role of A, you will see a "-".

=== *The following part was specific to the conditions in which second-order information was available (COSTLY and FREE). We use square brackets to indicate where instructions differed between the two [with information specific for the COSTLY treatment in brackets] ===*

Information

When you are taking your decision, you have two sorts of information at your disposal.

- 1. Your screen will display the three most recent decisions that the participant in role B has taken, when he was in the role of A. These three decisions of the participant in role B will always be displayed to you, and displayed in three boxes on your screen. These boxes are given in random order. When B has not yet been in the role of A, you will see a "-".*

2. *About one of these three decisions of the participant with role B, you can request additional information.*

If you select the option 'additional information' under one of the decisions, the information that B had, when he took this decision will be shown to you. This information consists of the three most recent decisions of the participant with which B was paired at that time. These decisions are also given in random order. You can confirm your decision by clicking the button 'Continue'. [Viewing this information is costly. These costs are 5 points.]

NB: you can request additional information for only one of the decisions of the participant with role B. If you select 'additional information' for more than one decision, an error message will appear on your screen.

In summary:

Always available: the most recent three decisions of participant B

Optionally available [at a cost of 5 points]: additional information about one decision of participant B, namely the three most recent decisions of the participant with which B was paired at that time.

=== the rest of the instructions were identical for each of the experimental conditions ===

The end of a round

When participant A has made his decision, both participants are informed about their score in this round, and their total amount of points at that moment. This is the end of the round. You will be paired to a new participant. Randomly, you will be assigned the role of A or B.

The end of the session

When the experiment has ended, a questionnaire will follow.

Please remain seated after filling out the questionnaire, until we call your desk number. You will be paid individually in the reception room.

Chapter eight

General discussion

The success of humans as a species depends to a large extent on our sociality. Two aspects of human nature are particularly unique: the unlevelled ability to cooperate in groups and the great capacity for social learning (Boyd et al. 2011; Whiten and Erdal 2012). By cooperating in groups, humans are able to surmount challenges impossible for any individual to deal with on its own. By learning from others, humans are able to rapidly adjust their behaviour to a vast array of different circumstances. The passing of information through social learning has been called ‘cultural transmission’ (Cavalli-Sforza and Feldman 1981). It has been argued that social learning gives rise to an inheritance system that is in many ways analogous to genetic transmission, and can lead to a process of ‘cultural evolution’ in which the frequencies and distributions of socially learned traits changes over time (Dawkins 1976; Cavalli-Sforza and Feldman 1981; Boyd and Richerson 1985).

The direction and outcome of cultural evolution depends on the social learning strategies that individuals use. Theoretical analyses suggest that due to the differences between cultural and genetic transmission, cultural evolution can more readily explain the unlevelled scale, diversity and degree of human cooperation (Boyd and Richerson 1985; Henrich and Boyd 1998; Richerson and Boyd 2004). As modelling studies on the cultural evolution of cooperation are accumulating, understanding of this process remains limited, largely because of a lack of systematic empirical evidence of how people actually learn from each other. In particular, almost nothing is known about human social learning in social contexts (*i.e.*, contexts of cooperation), where the outcome of behaviour depends on what others do. I set out to investigate the interplay of social learning and human cooperation, using an approach that combines developing cultural evolution theory with experiments contributing to its empirical basis.

This thesis

In this thesis, my investigation starts with the development of game theoretic models to map out how forms of social learning affect the

direction and outcome of cultural evolution (Chapters 2 and 3). In particular, Chapter 2 offers a systematic analysis of how conformism and payoff-based learning can drive evolutionary dynamics in a range of different social settings. The models I present here suggest that conformism can change basins of attraction of equilibria in coordination games and destroy coexistence equilibria in evasion games. In a social dilemma (such as the famous Prisoner's dilemma), conformism can stabilize a cooperative equilibrium. In contrast to influential views in the literature, this study suggests that conformism tends to hinder, rather than facilitate the evolution of cooperation by cultural group selection. Chapter 3 further discusses the potential roles of cultural group selection in the evolution of human cooperation. Using extensive individual-based simulations, I study to what extent different forms of group selection can promote cooperation, given that individuals use specific ways to learn from their peers. The analysis of this model indicates that the outcome of cultural evolution strongly depends on the interaction between the forms of social learning and group selection: although conformism can promote the evolution of cooperation when cooperative groups can replace non-cooperative groups, following a leader tends to promote cooperation under a wider range of conditions.

In Chapters 4 and 5 of this thesis, I aim to contribute to the experimental groundwork for an evidence-based theory of cultural evolution. By means of decision making experiments with human subjects, I explore what kinds of social information people use when confronted with a variety of social and non-social contexts. The experimental results presented in Chapter 4 show that people vary strongly in their social learning behaviour. Some people never consult others in deciding what to do, whereas others frequently rely on social learning in determining their behaviour. People differ even more strongly with respect to the kind of information they base their decisions on: some individuals mostly focus on payoffs of their peers, whereas others ignore these payoffs altogether and attend to the behaviour of the majority in their environment. Moreover, individual social learning strategies tend to be consistent over contexts: individuals who use a form of social learning in one social context tend to use the same form in other contexts. Further, I discuss how such individual variation may affect the course of cultural evolution.

In Chapter 5, I present experimental evidence for the time stability of individual differences in social learning strategies. In addition, this chapter explores the effects of assortment of social learning strategies on the dynamics of cooperation in a social dilemma: I find that groups of individuals that focus on the frequencies of behaviour in their social environment tend to achieve higher levels of cooperation than groups of individuals that focus on their peers' payoffs.

In Chapters 6 and 7, I consider the issue of how human cooperation can be supported by reciprocal helping. In many situations, people condition their cooperative behaviour on the past actions of their interaction partners. Chapter 6 presents the results of an experiment addressing how people use information from direct experience ("was this person nice to me before?") and reputations ("was this person nice to others before?") in decisions to help. These two kinds of social information are associated with direct and indirect reciprocity, respectively. I find that on average, information from direct experience tends to be more important in deciding to help someone. Furthermore, I observe that direct experience and reputations can compensate for one another; when personal experience with an interaction partner is bad, a good reputation can substantially increase helping rates. Conversely, however, when direct interactions are mostly positive, a bad reputation has only a minor effect. Chapter 7 presents the results of an experiment focusing on indirect reciprocity. I investigate here the extent to which people care about the motivations behind the actions of their interaction partners. Theory suggests that taking into account these motivations can efficiently stabilise cooperation in human groups, since 'justified' refusals to help defectors does not damage an individual's reputation. The experimental results indicate that people vary in the extent to which these motivations play a role in their decision making: some people merely reward those that tend to be helpful (I call this 'shallow reciprocity'), whereas others take into account the motivations behind the decisions of their interaction partners, and distinguish 'justified' from 'unjustified' helping behaviour (I call this 'deep reciprocity'). Chapter 7 ends with discussing the potential causes and consequences of this diversity of morals for the dynamics of human cooperation.

I conclude this final chapter by reflecting on the insights and limitations of the studies presented here. In addition, I will highlight a few issues concerning social learning and human cooperation, and sketch possible avenues for further theoretical and empirical research.

Abstractions and simplifications

Throughout this thesis I have simplified matters considerably, and abstracted away from various aspects of reality that likely influence social learning and its effects on human cooperation. The conceptual models of Chapters 2, 3 and 4 single out conformism, payoff-based learning, and leader-following, and analyse dynamics of cultural evolution using highly idealised versions of these specific forms of social learning in simple contexts of interaction. These abstractions and simplifications promote mathematical tractability and clarity of the analyses: they help understand the effects of forms of social learning on a very basic level, leaving out potentially confounding factors. Similarly, the decision making experiments presented in Chapters 4 to 7 consider relatively simple situations designed to reflect the essence of specific forms of human social interactions. Studying human behaviour in simple and stylised experimental situations facilitates comparisons to theoretical predictions (see, for instance, Figure 1 of Chapter 4, and Figure 1 of Chapter 5). Also, the controlled conditions of these experiments allow for exclusion of a range of confounding factors faced by studies that investigate human behaviour in a more ‘natural’ setting. However, using ‘cartoon versions’ of the real world in both models and experiments limits the explanatory power of these studies. Below I will delineate some abstractions and simplifications that may weaken my ability to draw strong conclusions from the research presented here.

In each of the studies presented in this thesis, individuals have only two behavioural options. Obviously, this is an overly simplistic representation of reality. In their day-to-day lives, people constantly have to choose courses of action from a vast array of behavioural options present. These options may differ in kind (“am I going for a hunt today, or am I going to collect berries”) or may reflect a continuum between two extremes (“how

much risk am I willing to take to protect my tribe from intruders from the outside”). In Chapters 2 and 3, the traits that determine behaviour are also binary: individuals have the tendency to choose one out of two behaviours, and this tendency can change over time through social learning. The dynamics of cultural evolution would likely be different when individuals would employ a mixed strategy specifying the probability to perform a certain behaviour, or a degree to which a behaviour is performed (*e.g.*, a level of cooperation). Processes of cultural transmission would take different forms when traits vary on a continuous scale. Conformism, for instance, would need to be modelled in a different way, since simply following the majority behaviour in a (small) sample of peers would not suffice.

Studying behaviour as making decisions between two discrete alternatives also limits my ability to draw strong conclusions from the experiments of Chapters 4 and 5. With regard to decision making itself, biases may arise when choice options have different degrees of ‘prominence’. For instance, it is plausible that individuals are inclined to choose options that are displayed at the top of the screen. Similarly, in collecting social information among their peers, the subject displayed as ‘Player 1’ may be consulted more frequently. In addition, the fact that the experimental design uses binary choices may affect individuals’ social information use in other ways. On the one hand, behaviour in the real world is less clear-cut and unambiguous than pulling a lever on a two-armed bandit. For instance, when gauging your social environment, the cooperativeness of your peers is often far from obvious (and estimating your peers’ payoffs may be even harder if not impossible). When the behaviour of others is ambiguous and not readily imitated through observation, non-social learning can be a more efficient way of finding out what behaviour is optimal. As a consequence, social learning might be less important than suggested by our experimental studies. On the other hand, the real world is obviously much more complex than the simple situations I consider here. For instance, when individuals have to learn to distinguish edible from poisonous items from a vast array of available foods, non-social learning is likely slow and costly and individuals may well resort to social learning. Because our experiments reflect more simple situations with fewer options, the degree to which individuals rely on social learning

in our experimental setting may be much lower than can be expected in reality.

One of the main artefacts of this laboratory setting is that interactions take place anonymously via a computer screen (see the screenshots in the Appendices of Chapters 4 and 7; sections 13 and 3, respectively). Participating subjects sit in cubicles in the laboratory and cannot look each other in the face while interacting. The exclusion of psychological effects of face-to-face encounters likely has considerable effects on the dynamics of social interactions. In many situations, people have to coordinate and cooperate in relatively small groups, and in such situations anonymity seems a rather strange condition. In day-to-day life, anonymous interactions are limited to quite specific contexts (such as e-trading or other social interactions on the internet). To me, laboratory studies of the kind I present in this thesis allow for getting valuable insights in basic aspects of human behaviour, and help develop general ideas on how social learning strategies might work. However, for the reasons given above, a limited picture of human behaviour emerges when one considers such studies in isolation. From this perspective, complementing elementary insights from the laboratory with field evidence would not only yield a more complete account for human behaviour, but also help demarcate the external validity of claims based on laboratory experiments.

The representation of social learning I employ in this thesis strongly limits the ways in which individuals can modify their behaviour, and particularly, and I omit individuals' past experience from consideration. In Chapters 2, 3 and 4 of this thesis, cultural traits are modelled as discrete units of inheritance that have exactly the same effect in different individuals. Individuals cannot refer to any past experience beyond their most recent behaviour and the payoff associated with it; adjustment to (social) environment only occurs by occasional copying of peers based on observation of behaviour, payoffs, frequencies or by simply following a leader in a group. A more comprehensive consideration of social learning would take into account how socially acquired information is integrated in the body of information present in individual experience. This integration may cause socially learned information to have different effects in different individuals. Within the specific context from which the

information is derived, individuals may gradually adjust their behavioural tendencies in the direction of the observed behaviour (for instance, cooperators may become slightly less cooperative by observing defectors). This integration may also cause changes in behavioural tendencies, or beliefs, values and ideas in other contexts, by generalisation and application of useful information in other situations (for instance, important insights from one domain of science can be often be fruitfully transported to another). Assessing the effects of such ‘spill-over’ processes present a perhaps difficult, but likely rewarding challenge for future models of cultural evolution.

Throughout this thesis I limit myself to highly stylised interactions, thereby ignoring the possibility that in real life the structure of the social context might be unclear to interaction partners. In each of the chapters I use ‘games’ with specific payoff matrices, aiming to capture essential aspects of social situations in which human regularly interact (*i.e.*, social dilemmas designed to mirror situations in which individual and group interests are opposed to each other; coordination games aiming to reflect situations in which individuals simultaneously have to settle on the same course of action; and evasion games in which individuals profit from performing behaviour that is rare). In my models, individuals can over time adjust their behaviour according to the payoff matrices underlying the games, and in each of the experiments presented here, participants are always informed about the general structure of the interaction contexts they are confronted with. In reality, however, it is typically unclear in what kind of social situation you find yourself – it is hard to imagine situations in which one can fathom the underlying payoff structure of a social situation, and condition behaviour on the material incentives to all individuals involved. In this light, it would be interesting to study social learning strategies in experiments in which individuals have limited information (or no information at all) about the underlying payoff structures of the contexts they interact in. Studying human behaviour in such ‘fuzzy’ contexts in the controlled environment of a decision making laboratory seems promising to get a more realistic picture of how people respond to the behaviour of others and may reveal interesting features of human social learning strategies.

Individual variation

As argued above, one must be careful with drawing conclusions about human behaviour from decision making experiments of the kind presented in this thesis. One aspect that is widely viewed as the most limiting factor in such studies is the fact that the subject pool is confined to subjects with a cultural background that can be described as western, educated, industrialized, rich and democratic ('WEIRD'; Henrich et al. 2010) – and the studies presented here are no exception. Over recent years, cross-cultural experiments have established that human social behaviour differs strongly between different cultures (Henrich et al. 2001; Herrmann et al. 2008; Gächter et al. 2010). From this perspective, it is likely that social learning strategies also vary between cultures (Efferson et al. 2007). As far as I know, no systematic empirical evidence exists on this topic as yet, but it seems plausible that the way in which social interactions in a population are structured influences the social learning strategies of its members. For instance, in societies built around cooperative social structures and in which people tend to generally trust each other, valuable information may be widely transmitted in social interactions. Accordingly, individuals may strongly rely on various forms of social learning. By contrast, in low-trust societies in which social information is scarce or unreliable (*e.g.*, due to manipulation) social learning may take other forms, or individuals may largely resort to non-social learning. In turn, population-level differences in social learning strategies likely lead to differences in the speed at which new technologies spread in these populations, but also the extent to which ideas, beliefs and social norms can gain a foothold.

The experiments presented in this thesis revealed strong individual differences in human social behaviour on a smaller scale, *within* a demographically homogeneous subject pool. Firstly, the observed individual differences (in social learning strategies and reciprocal strategies) in Chapters 4 to 7 illustrate that the *average* behaviour is often not representative for any single individual in the population. Furthermore, the models presented in Chapter 4 reveal that accounting for individual differences in social learning strategies matters for the

direction and outcome of cultural evolution. This reveals a potential weakness of all models of cultural evolution that I know of (including the ones presented in Chapters 2 and 3). Many of these models do assume that individuals use some mixture of social learning rules, such as payoff-based learning and conformism (*e.g.*, Henrich and Boyd 1998; Bowles 2001; Henrich 2001; Bowles et al. 2003), but theory has thus far ignored the implications of individual differences in social learning strategies. As a consequence, existent models likely give a misleading picture of how socially learned traits spread through populations, but also of how human cooperation may be shaped by cultural evolution.

The pronounced and individual differences that I find in each of the experimental studies presented here – and in particular, the finding that social learning strategies are consistent over time and contexts – makes one wonder whether they reflect deeper differences in how human behaviour is structured. The results of Chapter 5 of this thesis give support for the idea that the strategies measured in the laboratory reflect individuals' behavioural tendencies in the real world. I observe that at two different time points (with one month in between), individuals by and large use the same social learning behaviour, even when the description of the circumstances was changed. This suggests that such individual differences are not just some artefact arising from specific experimental conditions, but that people differ in the way their behaviour is structured – and take these behavioural tendencies with them into the laboratory.

Further investigation into the associations between social learning strategies and cooperative tendencies of individuals presents an interesting avenue of research for the future. For instance, payoff-based learning may be associated with selfish payoff-maximising behaviour, and may therefore be correlated with free-riding in social dilemmas (as suggested by the results presented in Chapter 5). By contrast, frequency-based learning may be associated with forms of 'conditional cooperation': individuals might keep track of frequencies of behaviour because they are only willing to pay the cost of cooperation when sufficient interaction partners also do so. Looking into correlations across contexts and investigating associations of traits that are involved in behaviour in social interaction will help understand how human behavioural tendencies are

structured, and to what extent individual behaviour is explained by distinct ‘personalities’.

In Chapter 7 I highlighted one obvious possible consequence of such consistent variation for human cooperation. Variation in moral standards (*i.e.*, reciprocal strategies) may well be correlated with other individual traits that function as signals based on which individuals select interaction partners. Such an assortment process may greatly affect the dynamics of cooperation: groups of deep reciprocators will likely maintain stable reputation-based cooperation, whereas cooperation is likely to break down among shallow reciprocators when they are confronted with defectors.

The nature of social learning strategies

Our understanding of cultural evolution would not only profit from having a solid empirical basis for social learning strategies, but also from a basic theory of how specific social learning strategies could have evolved in the first place. In Chapters 2 and 3 I simply *assume* that individuals use certain strategies of social learning, and ignore the issue of how forms of social learning strategies themselves evolve under the influence of selection. Theoretical analyses suggest that under a wide range of conditions, payoff-based learning and forms of conformism can be favoured by natural selection, in the sense that such strategies can invade a population of individuals that learn non-socially (*e.g.*, Boyd and Richerson 1985; Kendal et al. 2009; Rendell et al. 2010a). These analyses, however, typically consider only a small number of social learning strategies, and limit themselves to situations in which these strategies have to compete with strategies deploying non-social learning. This is strange, because populations likely harbour a variety of social learning strategies – which is illustrated by the diversity in social learning strategies observed in Chapters 4 and 5. Individuals using different social learning strategies likely differ in their efficiency to acquire useful information about their (social) environment, so that individuals with more effective social learning strategies will perform better. Based on this differential performance – reflected in *e.g.*, survival, payoffs, health, wealth or social

status – processes of selection can give rise to an evolutionary dynamic in which frequencies of alternative social learning strategies change over time, and over time more effective strategies tend to prevail.

To my knowledge, only one study explicitly pits a range of social learning strategies against each other (Rendell et al. 2010b). This study involved a tournament in which entrants (scientists as well as non-scientists) were invited to design a social learning strategy that would enable agents to survive in a simulated environment that changed over time, and over 100 submitted strategies were systematically subjected to competition between each other. This study gave interesting insights in when individuals should effectively copy others, and how to account for historical information when environments change from time to time. Although this tournament provides a valuable first step in studying the efficiency of various social learning strategies, this approach does not consider an explicit evolutionary process in which variation emerges by mutation, and differential transmission leads to gradual change over time. Rather, the strategies considered are pre-designed, and often make use of complex forms of information processing requiring massive computational power of individuals to actually deploy.

To gain insight in the evolution of social learning strategies, it would be informative to develop models in which social learning strategies emerge in a more gradual manner. Such an approach would not analyse differences in performance of a set of predefined strategies (with arbitrary implementations), but consider social learning strategies arising from scratch, whose aspects gradually change over time (*e.g.*, due to small mutations). Starting out from non-social learning, strategies could evolve towards relying to some extent on social learning, possibly conditioned on, and weighted with, individual experience. At the same time, strategies might develop a costly capacity to estimate the performance of peers (*e.g.*, reflected in their payoffs), optimise the way they react to frequencies of peers behaviour (*e.g.*, in a conformist or anti-conformist manner), or develop the quality of actively influencing peers' behaviour by presenting themselves as a leader or a teacher. Ideally, such models would also delineate the circumstances under which stable variation in social learning strategies (*cf.* Chapters 4 and 5 of this thesis) is a likely outcome of evolution.

Models aiming to explain the evolution of social learning strategies face a perhaps even more fundamental issue. Namely, nothing is known about how social learning strategies themselves transmit between individuals. Twin studies have shown that various human social traits, such as cooperative tendencies and preferences for fairness, have a genetic basis (Wallace et al. 2007; Cesarini et al. 2008). It thus seems plausible that there is some genetic basis to social learning strategies (so that these strategies could be subject to genetic evolution by natural selection). An interesting alternative possibility is that social learning strategies themselves are acquired through social learning. Parents have a strong influence on the behaviour of their children while raising them – and likely shape their offspring's social tendencies in this process. The way their offspring learns from others may be no exception. It may require clever experiments to disentangle genetic and cultural influences on the transmission of social learning strategies. Following the approach of the work mentioned above (Wallace et al. 2007; Cesarini et al. 2008), it seems promising to study the basis of social learning strategies by investigating social learning and cooperative behaviour of heterogametic and homogametic twins. A subsequent comparison of twins raised together and twins raised apart would lead to valuable insights in the genetic and cultural factors that constitute social learning strategies.

Concluding remarks

After four years of studying the interplay of social learning and human cooperation, I have to conclude that I only scratched the surface of these intriguing aspects of human nature. My models and experiments considered very simple social situations that reflected a rather coarse-grained picture of human interactions in the real world. A vast amount of work is still needed to bridge the gap between basic insights from the laboratory and the tremendous complexity of real-life social interactions. Perhaps an even bigger challenge lies in revealing the nature of social learning strategies, showing how such strategies are represented in individuals, how they transmit and how they evolve.

One comforting thought is that this field of research is relatively young. We are just getting started with investigating the intricate ways in which people acquire and use social information to modify their behaviour. Major issues lie ahead of us in understanding the processes underlying social learning and how they shape human behaviour in social and non-social situations. Despite the complexity of these matters, we have to take up the challenge to gain deeper insights into how social learning affects behaviour across the range of contexts in which humans interact with each other. Eventually, these insights into the interplay between social learning and cooperation will lead to a deeper understanding of what makes us humans unique: our sociality.

References

- Alexander R. D. (1987). *The Biology of Moral Systems*. Transaction Books, New York.
- Asch S. E. (1956). Studies of independence and conformity: I. A minority of one against a unanimous majority. *Psychological Monographs: General and Applied* **70**: 1-70.
- Axelrod R. (1984). *The Evolution Of Cooperation*. Basic Books, New York.
- Axelrod R., and W. D. Hamilton (1981). The evolution of cooperation. *Science* **211**: 1390-1396.
- Barrett H. C., L. Cosmides, and J. Tooby (2007). The hominid entry into the cognitive niche. In: S. W. Gangestad, and J. A. Simpson (eds.). *Evolution of Mind, Fundamental Questions and Controversies*. Guilford Press, New York. pp. 241-248.
- Bates D., M. Maechler, and B. Bolker (2012). lme4: Linear mixed-effects models using S4 classes.
- Battesti M., C. Moreno, D. Joly, and F. Mery (2012). Spread of social information and dynamics of social transmission within *Drosophila* Groups. *Current Biology* **22**: 309-313.
- Bolton G. E., E. Katok, and A. Ockenfels (2005). Cooperation among strangers with limited information about reputation. *Journal of Public Economics* **89**: 1457-1468.
- Botero C. A., I. Pen, J. Komdeur, and F. J. Weissing (2010). The evolution of individual variation in communication strategies. *Evolution* **64**: 3123-3133.
- Bowles S. (2001). Individual interactions, group conflicts, and the evolution of preferences. In: S. N. Durlauf, and H. Peyton Young (eds.). *Social Dynamics*. MIT Press, Cambridge. pp. 155-190.
- Bowles S. (2009). Did warfare among ancestral hunter-gatherers affect the evolution of human social behaviors? *Science* **324**: 1293-1298.
- Bowles S., J. Choi, and A. Hopfensitz (2003). The co-evolution of individual behaviors and social institutions. *Journal of Theoretical Biology* **223**:135-147.
- Boyd R., and P. J. Richerson (1985). *Culture and the Evolutionary Process*. University of Chicago Press, Chicago.

References

- Boyd R., and P. J. Richerson (1988). An evolutionary model of social learning: the effects of spatial and temporal variation. In: T. R. Zentall, and B. G. Galef (eds.). *Social Learning: Psychological and Biological Perspectives*. pp. 29-48.
- Boyd R., and P. J. Richerson (2002). Group beneficial norms can spread rapidly in a structured population. *Journal of Theoretical Biology* **215**: 287-296.
- Boyd R., H. Gintis, S. Bowles, and P. J. Richerson (2003). The evolution of altruistic punishment. *Proceedings of the National Academy of Sciences U. S. A.* **100**: 3531-3535.
- Boyd R., P. J. Richerson, and J. Henrich (2011a). The cultural niche: why social learning is essential for human adaptation. *Proceedings of the National Academy of Sciences U. S. A.* **108**: 10918-10925.
- Boyd R., P. J. Richerson, and J. Henrich (2011b). Rapid cultural adaptation can facilitate the evolution of large-scale cooperation. *Behavioral Ecology and Sociobiology* **65**: 431-444.
- Brandt H., and K. Sigmund (2004). The logic of reprobation: assessment and action rules for indirect reciprocation. *Journal of Theoretical Biology* **231**: 475-486.
- Brosig J., T. Riechmann, and J. Weimann (2007). Selfish in the end?: an investigation of consistency and stability of individual behavior. FEMM Working Paper Series 5. pp. 1-33.
- Cavalli-Sforza L. L., and M. W. Feldman (1981). *Cultural Transmission and Evolution: a Quantitative Approach*. Princeton University Press, Princeton.
- Cesarini D., C. T. Dawes, J. H. Fowler, M. Johannesson, P. Lichtenstein, and B. Wallace (2008). Heritability of cooperative behavior in the trust game. *Proceedings of the National Academy of Sciences U. S. A.* **105**: 3721-3726.
- Choi J., and S. Bowles (2007). The coevolution of parochial altruism and war. *Science* **318**:636-640.
- Clutton-Brock T. (2009). Cooperation between non-kin in animal societies. *Nature* **462**: 51-57.
- Darwin C. (1859). *The Origin of Species by Means of Natural Selection: or, the Preservation of Favoured Races in the Struggle for Life*. John Murray, London.
- Darwin C. (1871). *The Descent of Man*. John Murray, London.
- Dawkins R. (1976). *The Selfish Gene*. Oxford University Press, Oxford.

References

- Digman J. M. (1990). Personality structure: emergence of the five-factor model. *Annual Review of Psychology* **41**: 417-440.
- Duffy G. A., T. W. Pike, and K. N. Laland (2009). Size-dependent directed social learning in nine-spined sticklebacks. *Animal Behavior* **78**: 371-375.
- Dufwenberg M., U. Gneezy, W. Güth, and E. E. van Damme (2001). Direct versus indirect reciprocity: an experiment. *Homo Oeconomicus* **18**: 19-30.
- Dugatkin L. A., and J. J. Godin (1993). Female mate copying in the guppy *Poecilia reticulata*: age-dependent effects. *Behavioral Ecology* **4**: 289-292.
- Dugatkin L. A., and H. K. Reeve (1997). *Cooperation Among Animals: an Evolutionary Perspective*. Oxford University Press, Oxford.
- Efferson C., P. J. Richerson, R. McElreath, M. Lubell, E. Edsten, T. M. Waring, B. Paciotti, and W. Baum (2007). Learning, productivity, and noise: an experimental study of cultural transmission on the Bolivian Altiplano. *Evolution and Human Behavior* **28**: 11-17.
- Efferson C., R. Lalive, P. J. Richerson, R. McElreath, and M. Lubell (2008). Conformists and mavericks: the empirics of frequency-dependent cultural transmission. *Evolution and Human Behavior* **29**: 56-64.
- Engelmann D., and U. Fischbacher (2009). Indirect reciprocity and strategic reputation building in an experimental helping game. *Games and Economic Behavior* **67**: 399-407.
- Fehr E., and U. Fischbacher (2003). The nature of human altruism. *Nature* **425**: 785-791.
- Fischbacher U. (2007). z-Tree: Zurich toolbox for ready-made economic experiments. *Experimental Economics* **10**: 171-178.
- Fischbacher U., and S. Gächter (2010). Social preferences, beliefs, and the dynamics of free riding in public good experiments. *American Economic Review* **100**: 541-556.
- Gächter S., and B. Herrmann (2009). Reciprocity, culture and human cooperation: previous insights and a new cross-cultural experiment. *Philosophical Transactions of the Royal Society B: Biological Sciences* **364**: 791-806.
- Gächter S., B. Herrmann, and C. Thöni (2010). Culture and cooperation. *Philosophical Transactions of the Royal Society B: Biological Sciences* **365**: 2651-2661.
- Galef Jr B. G. (1976). Social transmission of acquired behaviour: a discussion of tradition and social learning in vertebrates. In: J. S.

References

- Rosenblatt, R. A. Hinde, E. Shaw, and C. Beer (eds.). *Advances in the Study of Behavior*. Academic Press, New York. pp. 77–100.
- Galef Jr B. G., and K. N. Laland (2005). Social learning in animals: empirical studies and theoretical models. *Bioscience* **55**: 489-499.
- García J., and J. C. J. M. Van den Bergh (2011). Evolution of parochial altruism by multilevel selection. *Evolution and Human Behavior* **32**: 277-287.
- Gargiulo F., and J. J. Ramasco (2012). Influence of opinion dynamics on the evolution of games. *PLoS One* **7**:e48916.
- Gigerenzer G., P. M. Todd, and ABC Research Group (1999). *Simple Heuristics that Make Us Smart*. Oxford University Press, Oxford.
- Gintis H. (2000). *Game Theory Evolving: a Problem-Centered Introduction to Modeling Strategic Interaction*. Princeton University Press, Princeton.
- Godin J., and L. A. Dugatkin (1996). Female mating preference for bold males in the guppy, *Poecilia reticulata*. *Proceedings of the National Academy of Sciences U. S. A.* **93**: 10262-10267.
- Gosling S. D. (2001). From mice to men: what can we learn about personality from animal research? *Psychological Bulletin* **127**: 45-86.
- Griffin A. S., S. A. West, and A. Buckling (2004). Cooperation and competition in pathogenic bacteria. *Nature* **430**: 1024-1027.
- Griskevicius V., N. J. Goldstein, C. R. Mortensen, R. B. Cialdini, and D. T. Kenrick (2006). Going along versus going alone: when fundamental motives facilitate strategic (non-)conformity. *Journal of Personality and Social Psychology* **91**:281-294.
- Guala F. (2012). Reciprocity: Weak or strong? What punishment experiments do (and do not) demonstrate. *Behavioral and Brain Sciences* **35**: 1–59.
- Guerinot M. L. (1994). Microbial iron transport. *Annual Reviews in Microbiology* **48**: 743-772.
- Gumbah (2002). In: Poelmo, slaaf van het zuiden (voor meisjes die van dieren houden). Hummelinck Stuurman Theaterbureau.
- Guzmán A. R., C. Rodríguez-Sickert, and R. Rowthorn (2007). When in Rome, do as the Romans do: the coevolution of altruistic punishment, conformist learning, and cooperation. *Evolution and Human Behavior* **28**: 112-117.
- Hamilton W. D. (1964). The genetical evolution of social behaviour I and II. *Journal of Theoretical Biology* **7**: 1-52.

References

- Henrich J. (2004). Cultural group selection, coevolutionary processes and large-scale cooperation. *Journal of Economic Behavior & Organization* **53**: 3-35.
- Henrich J., and R. Boyd (1998). The evolution of conformist transmission and the emergence of between-group differences. *Evolution and Human Behavior* **19**: 215-241.
- Henrich J. (2001). Cultural transmission and the diffusion of innovations: Adoption dynamics indicate that biased cultural transmission is the predominate force in behavioral change. *American Anthropologist* **103**: 992-1013.
- Henrich J., R. Boyd, S. Bowles, C. Camerer, E. Fehr, H. Gintis, and R. McElreath (2001). In search of homo economicus: behavioral experiments in 15 small-scale societies. *American Economic Review* **91**: 73-78.
- Henrich J., and R. Boyd (2001). Why people punish defectors. *Journal of Theoretical Biology* **208**: 79-89.
- Henrich J., and R. McElreath (2003). The evolution of cultural evolution. *Evolutionary Anthropology: Issues, News, and Reviews* **12**: 123-135.
- Henrich J., and F. J. Gil-White (2001). The evolution of prestige: freely conferred deference as a mechanism for enhancing the benefits of cultural transmission. *Evolution and Human Behavior* **22**: 165-196.
- Henrich J., S. J. Heine, and A. Norenzayan (2010). Beyond WEIRD: Towards a broad-based behavioral science. *Behavioral and Brain Sciences* **33**: 111-135.
- Henrich J., and N. Henrich (2010). The evolution of cultural adaptations: Fijian food taboos protect against dangerous marine toxins. *Proceedings of the Royal Society B: Biological Sciences* **277**: 3715-3724.
- Henrich J., and J. Broesch (2011). On the nature of cultural transmission networks: evidence from Fijian villages for adaptive learning biases. *Philosophical Transactions of the Royal Society B: Biological Sciences* **366**: 1139-1148.
- Herrmann B., C. Thöni, and S. Gächter (2008). Antisocial punishment across societies. *Science* **319**: 1362-1367.
- Heyes C. M. (1994). Social learning in animals: categories and mechanisms. *Biological Reviews* **69**: 207-231.
- Hoppitt W., and K. N. Laland (2013). *Social Learning: An Introduction to Mechanisms, Methods, and Models*. Princeton University Press, Princeton.

References

- Hoppitt W. J., G. R. Brown, R. Kendal, L. Rendell, A. Thornton, M. M. Webster, and K. N. Laland (2008). Lessons from animal teaching. *Trends in Ecology & Evolution* **23**: 486-493.
- Inglehart R., and W. E. Baker (2000). Modernization, cultural change, and the persistence of traditional values. *American Sociological Review* **65**: 19-51.
- Insko C. A., R. H. Smith, M. D. Alicke, J. Wade, and S. Taylor (1985). Conformity and Group Size The Concern with Being Right and the Concern with Being Liked. *Personality and Social Psychology Bulletin* **11**: 41-50.
- John O. P., R. W. Robinson, and L.A. Pervin (2010). Handbook of Personality: Theory and Research. Guilford Press, New York.
- Johnstone R. A., and A. Manica (2011). Evolution of personality differences in leadership. *Proceedings of the National Academy of Sciences U. S. A.* **108**: 8373-8378.
- Junikka J., L. Molleman, P. Van den Berg, F. J. Weissing, and M. Puurtinen (in preparation). Effects of assortment of cooperative strategies on the formation and maintenance of group norms.
- Karlin S., and H. E. Taylor (1975). A First Course in Stochastic Processes. Academic press, Waltham.
- Kendal J., L. Giraldeau, and K. Laland (2009). The evolution of social learning rules: payoff-biased and frequency-dependent biased transmission. *Journal of Theoretical Biology* **260**: 210-219.
- Klein R. G. (1989). The Human Career. University of Chicago Press, Chicago.
- Feldmann, M. W., K. Aoki, and J. Kumm (1996). Individual versus social learning: evolutionary analyses in a fluctuating environment. *Anthropological Sciences* **104**: 209-232.
- Kurzban R., and D. Houser (2005). Experiments investigating cooperative types in humans: A complement to evolutionary theory and simulations. *Proceedings of the National Academy of Sciences U. S. A.* **102**: 1803-1807.
- Laland K. N. (2004). Social learning strategies. *Learning & Behavior* **32**: 4-14.
- Lehmann L., and L. Keller (2006). The evolution of cooperation and altruism—a general framework and a classification of models. *Journal of Evolutionary Biology* **19**: 1365-1376.
- Lehmann L., N. Perrin, and F. Rousset (2006). Population demography and the evolution of helping behaviors. *Evolution* **60**: 1137-1151.

References

- Lehmann L., and M. W. Feldman (2008). The co-evolution of culturally inherited altruistic helping and cultural transmission under random group formation. *Theoretical Population Biology* **73**: 506-516.
- Lehmann L., M. W. Feldman, and K. R. Foster (2008a). Cultural transmission can inhibit the evolution of altruistic helping. *American Naturalist* **172**: 12-24.
- Lehmann L., K. R. Foster, E. Borenstein, and M. W. Feldman (2008b). Social and individual learning of helping in humans and other species. *Trends in Ecology & Evolution* **23**: 664-671.
- Leigh E. G. (1983). When does the good of the group override the advantage of the individual? *Proceedings of the National Academy of Sciences U. S. A.* **80**: 2985-2989.
- Leigh J. (2010). The group selection controversy. *Journal of Evolutionary Biology* **23**: 6-19.
- Leimar O., and P. Hammerstein (2001). Evolution of cooperation through indirect reciprocity. *Proceedings of the Royal Society of London B: Biological Sciences* **268**: 745-753.
- Mace R., and R. Sear (2005). Are humans cooperative breeders? In: E. Volland, A. Chasiotis, and W. Schiefenhoewel (eds.). *Grandmotherhood: the Evolutionary Significance of the Second Half of Female Life*. Rutgers University Press, New Brunswick. pp. 143-159.
- Maynard Smith J. (1964). Group selection and kin selection. *Nature* **201**: 1145-1147.
- McElreath R., A. V. Bell, C. Efferson, M. Lubell, P. J. Richerson, and T. Waring (2008). Beyond existence and aiming outside the laboratory: estimating frequency-dependent and pay-off-biased social learning strategies. *Philosophical Transactions of the Royal Society B: Biological Sciences* **363**: 3515-3528.
- McElreath R., M. Lubell, P. J. Richerson, T. M. Waring, W. Baum, E. Edsten, C. Efferson, and B. Paciotti. (2005). Applying evolutionary models to the laboratory study of social learning. *Evolution and Human Behavior* **26**: 483-508.
- McNamara J. M., and O. Leimar (2010). Variation and the response to variation as a basis for successful cooperation. *Philosophical Transactions of the Royal Society B: Biological Sciences* **365**: 2627-2633.
- McNamara J. M., and F. J. Weissing (2010). Evolutionary game theory. In: T. Székely, A. J. Moore, and J. Komdeur (eds.). *Social Behaviour: Genes, Ecology and Evolution*. Cambridge University Press, Cambridge. pp. 88-106.

References

- Mesoudi A., A. Whiten, and R. Dunbar (2006). A bias for social information in human cultural transmission. *British Journal of Psychology* **97**: 405-423.
- Mesoudi, A. (2008). An experimental simulation of the “copy-successful-individuals” cultural learning strategy: adaptive landscapes, producer–scrounger dynamics, and informational access costs. *Evolution and Human Behavior* **29**: 350-363.
- Mesoudi A., A. Whiten, A. Mesoudi, and A. Whiten (2008). The multiple roles of cultural transmission experiments in understanding human cultural evolution. *Philosophical Transactions of the Royal Society B: Biological Sciences* **363**: 3489-3501.
- Mesoudi A. (2011a). An experimental comparison of human social learning strategies: payoff-biased social learning is adaptive but underused. *Evolution and Human Behavior* **32**: 334-342.
- Mesoudi A. (2011b). *Cultural Evolution: How Darwinian Theory Can Explain Human Culture and Synthesize the Social Sciences*. University of Chicago Press, Chicago.
- Milinski M., D. Semmann, T. C. Bakker, and H. Krambeck (2001). Cooperation through indirect reciprocity: image scoring or standing strategy? *Proceedings of the Royal Society of London B: Biological Sciences* **268**: 2495-2501.
- Molleman L., P. van den Berg, and F. J. Weissing. Consistent individual differences in social learning strategies. (subm.).
- Molleman L., I. Pen, and F. J. Weissing (2013a). Effects of conformism on the cultural evolution of social behaviour. *PLoS One* **8**:e68153.
- Molleman L., A. E. Quiñones, and F. J. Weissing (2013b). Cultural evolution of cooperation: the interplay between forms of social learning and group selection. *Evolution and Human Behavior* **34**: 342-349.
- Molleman L., E. van den Broek, and M. Egas (2013c). Personal experience and reputation interact in human decisions to help reciprocally. *Proceedings of the Royal Society B: Biological Sciences* **280**: 20123044.
- Morgan T., L. Rendell, M. Ehn, W. Hoppitt, and K. Laland (2012). The evolutionary basis of human social learning. *Proceedings of the Royal Society B: Biological Sciences* **279**: 653-662.
- Noë R., and P. Hammerstein (1994). Biological markets: supply and demand determine the effect of partner choice in cooperation, mutualism and mating. *Behavioral Ecology and Sociobiology* **35**: 1-11.
- Nowak M. A. (2006). Five rules for the evolution of cooperation. *Science* **314**: 1560-1563.

References

- Nowak M. A., and K. Sigmund (1998). Evolution of indirect reciprocity by image scoring. *Nature* **393**: 573-577.
- Nowak M. A., and K. Sigmund (2005). Evolution of indirect reciprocity. *Nature* **437**: 1291-1298.
- Nowak M. A., A. Sasaki, C. Taylor, and D. Fudenberg (2004). Emergence of cooperation and evolutionary stability in finite populations. *Nature* **428**: 646-650.
- Ohtsuki H., and Y. Iwasa (2004). How should we define goodness?—reputation dynamics in indirect reciprocity. *Journal of Theoretical Biology* **231**: 107-120.
- Ohtsuki H., and Y. Iwasa (2006). The leading eight: social norms that can maintain cooperation by indirect reciprocity. *Journal of Theoretical Biology* **239**: 435-444.
- Okasha S. (2004). Multilevel selection and the partitioning of covariance: a comparison of three approaches. *Evolution* **58**: 486-494.
- Okasha S. (2006). *Evolution and the Levels of Selection*. Clarendon Press, Oxford.
- Osborne M. J., and A. Rubinstein (1994). *Course in Game Theory*. The MIT Press, Cambridge.
- Ostrom E., J. Walker, and R. Gardner (1992). Covenants with and without a sword: Self-governance is possible. *The American Political Science Review* **86**: 404-417.
- Otto S. P., and T. Day (2011). *A Biologist's Guide to Mathematical Modeling in Ecology and Evolution*. Princeton University Press, Princeton.
- Oyserman D., H. M. Coon, and M. Kemmelmeier (2002). Rethinking individualism and collectivism: evaluation of theoretical assumptions and meta-analyses. *Psychological Bulletin* **128**: 3-72.
- Panchanathan K. (2011). Two wrongs don't make a right: The initial viability of different assessment rules in the evolution of indirect reciprocity. *Journal of Theoretical Biology* **277**: 48-54.
- Panchanathan K., and R. Boyd (2003). A tale of two defectors: the importance of standing for evolution of indirect reciprocity. *Journal of Theoretical Biology* **224**: 115-126.
- Peña J., H. Volken, E. Pestelacci, and M. Tomassini (2009). Conformity hinders the evolution of cooperation on scale-free networks. *Physical Review E* **80**: 016110.
- Pike T. W., and K. N. Laland (2010). Conformist learning in nine-spined sticklebacks' foraging decisions. *Biology Letters* **6**: 466-468.

References

- Pinker S. (2010). The cognitive niche: Coevolution of intelligence, sociality, and language. *Proceedings of the National Academy of Sciences U. S. A.* **107**: 8993-8999.
- Price G. R. (1970). Selection and covariance. *Nature* **227**: 520-521.
- Queller D. C. (1992). Quantitative genetics, inclusive fitness, and group selection. *American Naturalist* **139**: 540-558.
- Rand D. G., S. Arbesman, and N. A. Christakis (2011). Dynamic social networks promote cooperation in experiments with humans. *Proceedings of the National Academy of Sciences U. S. A.* **108**: 19193-19198.
- Ratnieks F. L., and T. Wenseleers (2008). Altruism in insect societies and beyond: voluntary or enforced? *Trends in Ecology & Evolution* **23**: 45-52.
- Réale D., N. J. Dingemanse, A. J. Kazem, and J. Wright (2010). Evolutionary and ecological approaches to the study of personality. *Philosophical Transactions of the Royal Society B: Biological Sciences* **365**: 3937-3946.
- Rendell L., L. Fogarty, W. J. E. Hoppitt, T. J. H. Morgan, M. M. Webster, and K. N. Laland (2011). Cognitive culture: theoretical and empirical insights into social learning strategies. *Trends in Cognitive Sciences* **15**: 68-76.
- Rendell L., L. Fogarty, and K. N. Laland (2010a). Roger's paradox recast and resolved: population structure and the evolution of social learning strategies. *Evolution* **64**: 534-548.
- Rendell L., R. Boyd, D. Cownden, M. Enquist, K. Eriksson, M. W. Feldman, L. Fogarty, S. Ghirlanda, T. Lillicrap, and K. N. Laland (2010b). Why copy others? Insights from the social learning strategies tournament. *Science* **328**: 208-213.
- Richerson P. J., and R. Boyd (2004). *Not By Genes Alone: How Culture Transformed Human Evolution*. University of Chicago Press, Chicago.
- Roberts B. W., and W. F. DelVecchio (2000). The rank-order consistency of personality traits from childhood to old age: a quantitative review of longitudinal studies. *Psychological Bulletin* **126**: 3-25.
- Roberts G. (2008). Evolution of direct and indirect reciprocity. *Proceedings of the Royal Society B: Biological Sciences* **275**: 173-179.
- Rogers A. R. (1988). Does biology constrain culture? *American Anthropologist* **90**: 819-831.
- Rogers A. R. (1990). Group selection by selective emigration: the effects of migration and kin structure. *American Naturalist* **135**: 398-413.

References

- Scheuring I. (2009). Evolution of generous cooperative norms by cultural group selection. *Journal of Theoretical Biology* **257**: 397-407.
- Schlag K. H. (1998). Why imitate, and if so, how?: a boundedly rational approach to multi-armed bandits. *Journal of Economic Theory* **78**: 130-156.
- Seinen I., and A. Schram (2006). Social status and group norms: Indirect reciprocity in a repeated helping experiment. *European Economic Review* **50**: 581-602.
- Sherman P. W. (1977). Nepotism and the evolution of alarm calls. *Science* **197**: 1246-1253.
- Sih A., A. Bell, and J. C. Johnson (2004). Behavioral syndromes: an ecological and evolutionary overview. *Trends in Ecology & Evolution* **19**: 372-378.
- Simon B., J. A. Fletcher, and M. Doebeli (2012). Towards a general theory of group selection. *Evolution* **67**: 1561-1572.
- Skyrms B. 2005. Dynamics of conformist bias. *The Monist* **88**: 260-269.
- Soltis J., R. Boyd, and P. J. Richerson (1995). Can group-functional behaviors evolve by cultural group selection?: An empirical test. *Current Anthropology* **36**: 473-494.
- Sommerfeld R. D., H. Krambeck, and M. Milinski (2008). Multiple gossip statements and their effect on reputation and trustworthiness. *Proceedings of the Royal Society B: Biological Sciences* **275**: 2529-2536.
- Sommerfeld R. D., H. Krambeck, D. Semmann, and M. Milinski (2007). Gossip as an alternative for direct observation in games of indirect reciprocity. *Proceedings of the National Academy of Sciences U. S. A.* **104**: 17435-17440.
- Stiner M. C., R. Barkai, and A. Gopher (2009). Cooperative hunting and meat sharing 400-200 kya at Qesem Cave, Israel. *Proceedings of the National Academy of Sciences U. S. A.* **106**:13207-13212.
- Sugden R. (1986). *The Economics of Rights, Co-operation and Welfare*. Basil Blackwell, Oxford.
- Toelch U., M. J. van Delft, M. J. Bruce, R. Donders, M. T. Meeus, and S. M. Reader (2009). Decreased environmental variability induces a bias for social information use in humans. *Evolution and Human Behavior* **30**: 32-40.
- Traulsen A., and M. A. Nowak. (2006). Evolution of cooperation by multilevel selection. *Proceedings of the National Academy of Sciences U. S. A.* **103**: 10952-10955.

References

- Traulsen A., M. A. Nowak, and J. M. Pacheco (2006). Stochastic dynamics of invasion and fixation. *Physical Review E* **74**: 011909.
- Traulsen A., D. Semmann, R. D. Sommerfeld, H. J. Krambeck, and M. Milinski (2010). Human strategy updating in evolutionary games. *Proceedings of the National Academy of Sciences U. S. A.* **107**: 2962-2966.
- Trivers R. L. (1971). The evolution of reciprocal altruism. *Quarterly Review of Biology* **46**: 35-57.
- Ule A., A. Schram, A. Riedl, and T. N. Cason (2009). Indirect punishment and generosity toward strangers. *Science* **326**: 1701-1704.
- Valone T. J. (2007). From eavesdropping on performance to copying the behavior of others: a review of public information use. *Behavioral Ecology and Sociobiology* **62**: 1-14.
- Van de Waal E., C. Borgeaud, and A. Whiten (2013). Potent social learning and conformity shape a wild primate's foraging decisions. *Science* **340**: 483-485.
- Van den Berg P., L. Molleman, and F. Weissing (2012). The social costs of punishment. Reply to: Guala, F. Reciprocity: Weak or strong? What punishment experiments do (and do not) demonstrate. *Behavioral and Brain Sciences* **35**. pp. 42-43.
- Venables W. N., D. M. Smith, and the R development core team (2005). An Introduction to R: A Programming Environment for Data Analysis and Graphics.
- Vilone D., J. J. Ramasco, A. Sánchez, and M. San Miguel (2012). Social and strategic imitation: the way to consensus. *Scientific Reports* **2**: 686.
- Volk S., C. Thöni, and W. Ruigrok (2012). Temporal stability and psychological foundations of cooperation preferences. *Journal of Economic Behavior & Organization* **81**: 664-676.
- Wakano J. Y., K. Aoki, and M. W. Feldman (2004). Evolution of social learning: a mathematical analysis. *Theoretical Population Biology* **66**: 249-258.
- Wallace B., D. Cesarini, P. Lichtenstein, and M. Johannesson. (2007). Heritability of ultimatum game responder behavior. *Proceedings of the National Academy of Sciences U. S. A.* **104**: 15631-15634.
- Wedekind C., and M. Milinski (2000). Cooperation through image scoring in humans. *Science* **288**: 850-852.
- Wedekind C., and M. Milinski (1996). Human cooperation in the simultaneous and the alternating Prisoner's Dilemma: Pavlov versus

References

- Generous Tit-for-Tat. *Proceedings of the National Academy of Sciences U. S. A.* **93**: 2686-2689.
- Weissing F. J. (2011). Animal behaviour: Born leaders. *Nature* **474**: 288-289.
- West S. A., A. S. Griffin, and A. Gardner (2007). Social semantics: altruism, cooperation, mutualism, strong reciprocity and group selection. *Journal of Evolutionary Biology* **20**: 415-432.
- West S. A., A. S. Griffin, and A. Gardner (2008). Social semantics: how useful has group selection been? *Journal of Evolutionary Biology* **21**: 374-385.
- Whiten A., and D. Erdal (2012). The human socio-cognitive niche and its evolutionary origins. *Philosophical Transactions of the Royal Society B: Biological Sciences* **367**: 2119-2129.
- Wilson D. S. (1975). A theory of group selection. *Proceedings of the National Academy of Sciences U. S. A.* **72**: 143-146.
- Wilson D. S., and E. O. Wilson (2007). Rethinking the theoretical foundation of sociobiology. *Quarterly Review of Biology* **82**: 327-348.
- Wilson D. (2008). Social semantics: toward a genuine pluralism in the study of social behaviour. *Journal of Evolutionary Biology* **21**: 368-373.
- Wit J. (1999). Social learning in a common interest voting game. *Games and Economic Behavior* **26**: 131-156.
- Wolf M., and F. J. Weissing (2012). Animal personalities: consequences for ecology and evolution. *Trends in Ecology & Evolution* **27**: 452-461.
- Wolf M., G. S. Van Doorn, and F. J. Weissing (2008). Evolutionary emergence of responsive and unresponsive personalities. *Proceedings of the National Academy of Sciences U. S. A.* **105**: 15825-15830.
- Wolf M., G. S. van Doorn, O. Leimar, and F. J. Weissing (2007). Life-history trade-offs favour the evolution of animal personalities. *Nature* **447**: 581-584.
- Wolf M., and F. J. Weissing (2010). An explanatory framework for adaptive personality differences. *Philosophical Transactions of the Royal Society B: Biological Sciences* **365**: 3959-3968.
- Wolf M., G. S. Van Doorn, and F. J. Weissing (2011). On the coevolution of social responsiveness and behavioural consistency. *Proceedings of the Royal Society B: Biological Sciences U. S. A.* **278**: 440-448.
- Yamagishi T. (1986). The provision of a sanctioning system as a public good. *Journal of Personality and Social Psychology* **51**: 110-116.

Dutch summary

Wat maakt de mens zo'n succesvolle soort? Welke factoren zorgen ervoor dat mensen kunnen overleven onder de meest uiteenlopende omstandigheden op aarde? Een verklaring hiervoor schuilt in het feit dat de mens een buitengewoon sociaal dier is. Twee sociale eigenschappen springen bijzonder in het oog. Ten eerste zijn mensen goed in staat om samen te werken in groepen, veel beter dan onze naaste verwanten, de gorilla en de chimpansee. Door efficiënte coördinatie van taken en effectieve samenwerking zijn mensen in staat om het hoofd te bieden aan uitdagingen, waartegen geen enkel individu in zijn eentje opgewassen is. Ten tweede hebben mensen een bijzondere aanleg voor sociaal leren: door te observeren hoe anderen omgaan met specifieke situaties kunnen mensen hun gedrag snel aanpassen aan een groot aantal uiteenlopende omstandigheden. In dit proefschrift onderzoek ik het samenspel tussen samenwerking en sociaal leren en kijk ik hoe deze twee eigenschappen bijdragen aan het succes van de mens. Voordat ik het werk uit mijn proefschrift beschrijf, is een nadere toelichting van de concepten samenwerking en sociaal leren op zijn plaats.

Samenwerking tussen mensen doet zich voor in veel verschillende vormen. Geen enkel ander dier lukt het om acties in een groep op elkaar af te stemmen in zo'n breed scala aan verschillende situaties; van het helpen bij de opvoeding van elkaars kinderen tot de gecoördineerde jacht op grote prooidieren, en van het verzorgen van de zieken tot het verdedigen van een dorp tegen bedreigingen van buitenaf. Deze samenwerking is veelal gebaseerd op het feit dat mensen bereid zijn iets op te geven voor anderen in hun omgeving. Deze bereidheid om te investeren in samenwerking met anderen levert voordelen op die ervoor gezorgd hebben dat groepen mensen zich konden aanpassen aan bijna alle leefomgevingen op aarde. Echter, vanuit een evolutionair perspectief is onze aanleg tot samenwerking niet vanzelfsprekend. Charles Darwin realiseerde zich al dat investeren in samenwerking vaak problemen met zich meebrengt. Namelijk, samenwerking kan vaak worden geëxploiteerd door 'free-riders' die profiteren van de investeringen van anderen zonder

zelf de kosten te dragen. Met andere woorden: natuurlijke selectie bevoordeelt free-riders ondermijnt samenwerking. Sinds de tijd van Darwin hebben evolutiebiologen tal van mechanismen onder de loep genomen die de evolutie van samenwerking zouden kunnen verklaren, en laten zien dat free-riding op de korte termijn vaak voordelig is, maar op langere termijn juist nadelig kan zijn. Een bekend voorbeeld van een mechanisme dat samenwerking kan bevoordelen is 'groepsselectie'. Dit treedt op als selectie niet alleen plaatsvindt tussen individuen, maar ook op het niveau van de groep. Samenwerking kan er bijvoorbeeld voor zorgen dat een groep grotere kansen heeft te overleven of niet uiteen te vallen, of om conflicten met andere groepen te winnen. Hierdoor kan samenwerking bevoordeeld worden boven free-riding. Ook kunnen verwante individuen elkaar helpen zodat hun gemeenschappelijke genen worden doorgegeven aan volgende generaties. Een derde manier waarop samenwerking kan lonen, is gebaseerd op wederkerigheid: partners kunnen profiteren als ze beurtelings iets voor elkaar over hebben.

Sociaal leren stelt individuen in staat hun gedrag aan te passen door observatie van – en interacties met – anderen. Dit kan veel voordelen hebben: door anderen te observeren en hun gedrag te kopiëren, wordt het mogelijk om waardevolle informatie te verkrijgen over van allerlei aspecten van de omgeving, zoals de locatie van voedsel, waar natuurlijke vijanden zich ophouden, en hoe je met anderen dient om te gaan. Evolutionaire modellen suggereren dat het vaak loont om op een strategische manier om te gaan met de informatie uit je sociale omgeving. Een 'sociale leerstrategie' wordt over het algemeen gemodelleerd als een set regels die beschrijft hoe een individu van anderen leert. Deel van zo'n strategie kan zijn dat een individu zijn gedrag baseert op het gedrag van anderen, als hij zelf niet zeker is wat hij moet doen. Ook kan zo'n strategie beschrijven *hoe* van anderen geleerd wordt; bijvoorbeeld door te conformeren aan de meerderheid in de sociale omgeving, of door het gedrag te kopiëren van een succesvol of invloedrijk individu. Sinds de jaren 1970 hebben auteurs als Richard Dawkins betoogd dat sociaal leren aan de basis ligt van een systeem van overerving dat op veel manieren lijkt op genetische overerving. Mensen geven namelijk niet alleen hun genen door, maar ook ideeën, meningen of overtuigingen hun gedrag beïnvloeden. Net als genen, kunnen deze 'memen' variëren binnen een

populatie, en deze variatie kan ontstaan als individuen bijvoorbeeld experimenteren met nieuw gedrag, of een verbeterde techniek ontdekken. Zulke 'innovaties' dragen, net zoals mutaties in het genetische systeem, bij aan de variatie van ideeën in de populatie. Bovendien verspreiden sommige ideeën zich gemakkelijker dan andere, met name als ze geassocieerd zijn met gedrag dat leidt tot succesvolle aanpassing aan de heersende omstandigheden, of gedrag dat leidt tot een hoge sociale status. Zulk 'payoff-gebaseerd' leren doet sterk denken aan natuurlijke selectie. Net als natuurlijke selectie, leidt payoff-gebaseerd leren vaak tot de verspreiding van gedrag dat is toegesneden op de vereisten van de omgeving. Het besef dat culturele en genetische overerving grote overeenkomsten vertonen, leidde tot de ontwikkeling van een theorie van 'culturele evolutie'. Deze theorie bouwt voort op inzichten uit genetische studies en kijkt naar variatie, selectie en overerving, om te begrijpen hoe gedrag evolueert over de tijd, onder invloed van sociaal leren.

Wat hebben samenwerking en sociaal leren met elkaar te maken? Er is gesuggereerd dat sociaal leren de sleutel is tot de evolutie van samenwerking bij mensen. Een invloedrijke reeks modellen ondersteunt het argument dat vergeleken met genetische evolutie, culturele evolutie – onder invloed van sociaal leren – een betere verklaring kan geven voor de diversiteit en het niveau van menselijke samenwerking. Deze modellen gaan ervan uit dat conformistisch leren een grote rol speelt bij het in stand houden van samenwerking in groepen. Vervolgens kan, volgens deze modellen, samenwerking zich verspreiden door een proces dat 'culturele groepsselectie' wordt genoemd: in competitie tussen groepen (bijvoorbeeld door een conflict) zijn coöperatieve groepen in het voordeel, en door deze selectie op het niveau van groepen, kan free-riding worden onderdrukt. Hieronder zal ik uiteenzetten hoe dit argument in elkaar steekt, en de logica ervan tegen het licht houden met verschillende modellen van culturele evolutie.

In dit proefschrift gebruik ik een combinatie van experimenten en theoretische modellen om het samenspel van sociaal leren en samenwerking te onderzoeken. In twee experimentele studies (hoofdstukken 6 en 7) bekijk ik welke strategieën mensen toepassen in een samenwerkings situatie. Met behulp van modellen (hoofdstukken 2 en 3) onderzoek ik hoe de culturele evolutie van coöperatief gedrag wordt

beïnvloed door sociale leerstrategieën en vormen van culturele groepsselectie. Met nog eens twee experimenten (hoofdstukken 4 en 5) bekijk ik op welke manieren mensen van elkaar leren, om bij te dragen aan een empirische grondslag voor de theorie van culturele evolutie. Ik gebruik hierbij modellen om de gevolgen in kaart te brengen van de gevonden leerstrategieën.

Dit proefschrift

Ik begin dit proefschrift met het in kaart brengen hoe de dynamiek van culturele evolutie wordt beïnvloed door verschillende vormen van sociaal leren. Hiertoe ontwerp ik een aantal modellen die culturele evolutie van samenwerking nabootsen die in kaart brengen hoe de dynamiek van culturele evolutie wordt beïnvloed door verschillende leerstrategieën. Als startpunt van alle modellen neem ik payoff-gebaseerd leren, en bekijk hoe de dynamiek van culturele evolutie wordt beïnvloed als andere vormen van sociaal leren – waaronder conformisme – een rol gaan spelen.

Zoals boven gezegd, leidt payoff-gebaseerd leren tot een dynamiek van culturele evolutie die lijkt op genetische evolutie onder natuurlijke selectie. Groepen waarin wordt samengewerkt lopen hierbij altijd het risico om uitgebuit te worden door free-riders, omdat deze een hogere payoff behalen. Hierdoor wordt samenwerking snel ondermijnd, en de meeste modellen – net als die van genetische evolutie onder natuurlijke selectie – voorspellen dat groepen dan ook snel zullen worden overgenomen door free-riders. Echter, payoff-gebaseerd leren is niet de enige sociale leerstrategie die mensen gebruiken; vaak hebben mensen een neiging om zich te conformeren aan de meerderheid in hun sociale omgeving. Met behulp van een reeks theoretische modellen hebben verschillende onderzoekers betoogd dat juist dit conformisme aan de basis ligt van de culturele evolutie van menselijke samenwerking. Conformisme draagt er namelijk aan bij dat in sommige groepen samenwerking kan stabiliseren, en free-riders buiten de deur worden gehouden doordat individuen zich conformeren aan de meerderheid, en goed blijven samenwerken. Vervolgens kan deze coöperatieve groep een groot voordeel hebben in competitie met andere groepen: de voorvechters

van dit argument beweren dat zo'n proces van culturele groepsselectie van grote invloed is geweest op de evolutie van menselijke samenwerking. In **hoofdstukken 2 en 3** van dit proefschrift houd ik deze argumentatie op basis van theorie van culturele evolutie tegen het licht.

In **hoofdstuk 2** analyseer ik culturele evolutie in samenwerkingsituaties die dagelijkse sociale interacties van mensen weerspiegelen. Ik maak gebruik van het 'sociale dilemma', een bekend model waarin de belangen van een individu tegengesteld zijn aan de belangen van zijn groep: een investering in samenwerking levert het beste resultaat op voor alle betrokkenen, maar de verleiding bestaat om te profiteren van andermans investeringen zonder zelf bij te dragen. Ik stel hierbij de vraag hoe de dynamiek van culturele evolutie verandert als conformisme een rol gaat spelen in het leergedrag van individuen, die in principe payoff-gebaseerd leren. Dit payoff-gebaseerd leren leidt over het algemeen tot ondermijning van samenwerking, omdat binnen een groep free-riders een hogere payoff halen door de samenwerking van anderen te exploiteren. Conformisme, van de andere kant, zorgt ervoor dat individuen zich gaan schikken naar gedrag van de meerderheid. In de literatuur wordt sterk benadrukt dat conformisme ervoor zorgt dat het idee van 'free-riding' in groepen niet makkelijk van de grond kan komen in groepen die sterk samenwerken. Op deze manier zorgt conformisme ervoor dat samenwerking in een coöperatieve groep stabiel zou kunnen blijven. Mijn model laat echter zien dat het omgekeerde ook geldt: het idee van 'coöpereren' kan in groepen van free-riders niet makkelijk van de grond komen. In groepen van free-riders werkt zowel payoff-gebaseerd leren als conformisme tot de ondermijning van samenwerking, want free-riders halen een hogere payoff, en zijn ook nog eens in de meerderheid. Het initiëren van samenwerking in een groep free-riders blijkt daardoor praktisch onmogelijk. Dit effect zorgt ervoor dat conformisme – in tegenstelling tot conclusies van bestaande theorieën gebaseerd op culturele groepsselectie – de evolutie van samenwerking lang niet altijd bevordert, en vaak juist bemoeilijkt. Dit inzicht draagt bij aan een breed gevoerde discussie in de literatuur over de rol van conformisme in de culturele evolutie van samenwerking.

Verder onderzoek ik met dit model de effecten van conformisme in twee andere contexten waarin mensen vaak samenwerken. In een

‘coördinatiespel’ (waarbij het voordelig is als individuen hun gedrag op elkaar afstemmen), kan conformisme de richting van culturele evolutie veranderen, vooral als er weinig overeenstemming in een groep bestaat over welk gedrag nu moet worden aangenomen. In een ‘anti-coördinatiespel’ (waarbij het voordelig is om juist af te wijken van de meerderheid) kan conformisme ook grote effecten hebben: als individuen leren op basis van payoffs, kan er een stabiel patroon ontstaan waarbij verschillende individuen verschillende rollen vervullen. Wanneer individuen echter leren op basis van conformisme, dan kan dit patroon worden verstoord: conformistisch leren leidt namelijk tot homogeniteit, en kan er voor zorgen dat diversiteit aan gedrag snel verloren gaat. In een situatie waarin het afwijken van de meerderheid voordelig is, kan conformisme de uitkomst van interacties voor alle individuen verslechteren.

De modellen uit hoofdstuk 2 zijn relatief eenvoudig opgezet. Dit maakt aan de ene kant, dat ze voor een groot deel wiskundig geanalyseerd kunnen worden. Van de andere kant blijven er, door de simplificaties die daarvoor gemaakt moeten worden, een aantal aspecten uit de realiteit onderbelicht. In **hoofdstuk 3** ontwerp ik een simulatiemodel dat meer realistisch is, om de invloed van culturele groepsselectie op de evolutie van samenwerking te belichten. Ik bekijk verschillende vormen van culturele groepsselectie die in de literatuur worden beschreven. Uit het simulatiemodel blijkt dat de manier waarop competitie tussen groepen precies plaatsvindt, essentieel is voor de uitkomst van culturele evolutie. Als coöperatieve groepen een grotere invloed hebben in de hele populatie, en hun ideeën zich makkelijker verspreiden naar andere groepen, heeft conformisme een negatief effect op de evolutie van samenwerking. Hoewel voor deze vorm van culturele groepsselectie enkele groepen een stabiele samenwerking kunnen bereiken, verspreidt samenwerking zich maar moeilijk naar andere groepen. Door conformisme krijgen nieuwe ideeën in een groep nauwelijks een voet aan de grond, en zullen groepen van free-riders onveranderd blijven. Dit bevestigt de conclusies uit hoofdstuk 2. Conformisme kan de evolutie van samenwerking alleen bevorderen als coöperatieve groepen een grotere kans hebben een conflict met andere groepen te winnen, en vervolgens hun ‘cultuur’ opleggen aan verliezende groepen. Over het algemeen kan samenwerking gemakkelijker evolueren

als individuen geneigd zijn een leider in een groep te volgen, in plaats van te conformeren aan een meerderheid.

Samengevat laten de modellen uit de hoofdstukken 2 en 3 zien dat de manier waarop individuen van elkaar leren essentieel is voor de uitkomst van culturele evolutie. Culturele groepsselectie leidt lang niet altijd tot de evolutie van samenwerking; de richting en uitkomst van culturele evolutie is afhankelijk van het samenspel tussen vormen van sociaal leren en de manier waarop groepscompetitie plaatsvindt.

Steeds meer modelmatig onderzoek richt zich op het begrijpen van hoe verschillende leerstrategieën het verloop van culturele evolutie beïnvloeden. Echter, ons begrip van het proces van culturele evolutie blijft beperkt. Dit komt vooral doordat er tot dusver niet veel systematisch empirisch onderzoek is gedaan naar hoe mensen in de praktijk van elkaar leren. In de **hoofdstukken 4 en 5** neem ik een aantal stappen om deze gaten in onze kennis op te vullen, en draag daarmee bij aan een theorie van culturele evolutie met een empirische grondslag. Hiervoor heb ik enkele experimenten uitgevoerd, waarbij het doel is om sociale leerstrategieën van mensen te bestuderen in verschillende sociale situaties. Hiervoor gebruikte ik methodes uit de experimentele economie, waarbij proefpersonen krijgen betaald voor hun deelname. Hoeveel ze precies verdienen, hangt af van hun eigen beslissingen en de beslissingen van anderen. In een computerlaboratorium worden mensen in verschillende groepen ingedeeld, en communiceren alleen met elkaar via hun scherm. De groepen worden geconfronteerd met verschillende sociale situaties (zoals een sociaal dilemma, een coördinatiespel en een anti-coördinatiespel, zoals die uit hoofdstuk 2), en speciale software houdt bij welke beslissingen de deelnemers maken. Het doel van de experimenten is om in beeld te brengen hoe mensen sociale informatie gebruiken om deze beslissingen te maken. Voordat ze een beslissing maken, kunnen deelnemers informatie verzamelen over de voorgaande beslissing van de anderen in hun groep. Deelnemers kunnen er bijvoorbeeld voor kiezen om de beslissingen van een aantal groepsleden te bekijken (hierdoor wordt bijvoorbeeld conformisme mogelijk), of om ook de opbrengst van deze beslissing te inspecteren (hierdoor wordt payoff-gebaseerd leren mogelijk).

In **hoofdstuk 4** laat ik zien dat mensen sterk van elkaar verschillen in hun sociaal leergedrag. Sommigen laten zich in het maken van hun beslissingen vooral leiden door hun sociale omgeving, terwijl anderen zich hier weinig tot niets van aantrekken, en op zichzelf vertrouwen. Bovendien zijn er grote verschillen tussen de mensen die zich laten leiden door hun sociale omgeving: waar sommige mensen zich concentreren op welk gedrag tot de hoogste payoffs leidt, negeren anderen dit soort informatie in het geheel, en hebben slechts aandacht voor wat de meerderheid doet. Dit is een verrassende uitkomst: theorieën van culturele evolutie gaan er doorgaans vanuit dat ieder individu in een populatie op dezelfde manier leert. Nog verrassender is de uitkomst als we kijken naar leergedrag van individuen in de verschillende sociale situaties. De resultaten laten zien dat mensen niet zo flexibel zijn in hun sociaal leergedrag als wel wordt gedacht, maar eerder vuistregels gebruiken om te bepalen wat ze doen. Geconfronteerd met zeer uiteenlopende situaties, zijn onverwacht veel mensen consistent in hun sociale leerstrategie: óf zij richten zich op het gedrag van de meerderheid, óf zij leren op basis van payoff. Om de gevolgen van deze individuele verschillen in sociale leerstrategieën voor culturele evolutie te onderzoeken, ontwikkel ik in dit hoofdstuk ook verschillende modellen. Hierin vergelijk ik heterogene groepen (waarin sommigen altijd payoff-gebaseerd leren, en anderen altijd op basis van wat de meerderheid doet) met homogene groepen (waarin alle individuen leren door een mix van deze twee vormen van sociaal leren). Deze modellen laten zien dat in sommige omstandigheden – zoals het anti-coördinatiespel – groepen met een diversiteit aan leerstrategieën beter kunnen functioneren, doordat zij op een efficiënte manier een patroon kunnen bewerkstelligen waarin verschillende individuen verschillende rollen vervullen.

In **hoofdstuk 5** onderzoek ik met een soortgelijk experiment of sociale leerstrategieën van individuen stabiel zijn over de tijd. Hiertoe meet ik de sociale leerstrategieën van proefpersonen op twee verschillende momenten (met vier weken ertussen), in twee verschillende experimenten. De vergelijking van de twee metingen laat zien dat leerstrategieën relatief stabiel zijn over de tijd. Dit ondersteunt het idee dat metingen in mijn experimenten gedrag uit het dagelijks leven – buiten het laboratorium – weerspiegelen, en dat zij niet slechts weergeven hoe

mensen zich toevallig in één experiment gedragen. Verder bekijk ik met dit experiment of leerstrategieën van invloed zijn op samenwerking in een groep. Dit blijkt het geval te zijn: groepen van individuen die zich richten op het gedrag van de meerderheid, bereiken een hoge mate van samenwerking, vergeleken met groepen van individuen die vooral aandacht hebben voor de payoffs van groepsgenoten.

Samengevat laten de experimenten van hoofdstuk 4 en 5 zien dat mensen sterk verschillen in hun sociale leerstrategieën. Deze leerstrategieën zijn consistent over verschillende sociale situaties en stabiel over tijd. Ook hangen ze samen met een neiging tot samenwerking. Dit doet vermoeden dat deze strategieën onderdeel zijn van een persoonlijkheidsstructuur, en zouden gerelateerd kunnen zijn aan andere karaktereigenschappen zoals die wel in de psychologie worden onderzocht. In **box 5.1** doe ik een eerste stap om dit te onderzoeken, en laat ik zien dat competitieve individuen minder geneigd zijn om met anderen samen te werken.

Hoofdstukken 6 en 7 gaan niet over sociaal leren, maar over de vraag welke gedragsstrategieën mensen toepassen in een samenwerkings situatie. Ik laat hier de resultaten zien van een set experimenten die inzicht geven in de toepasbaarheid van het idee dat samenwerking tussen mensen kan worden bevorderd door wederkerig helpen. In theorieën van samenwerking worden twee vormen van wederkerigheid intensief onderzocht: *i*) 'directe wederkerigheid', waarin directe ervaringen uit voorgaande interacties bepalen of iemand bereid is een ander te helpen; en *ii*) 'indirecte wederkerigheid', waarin de reputatie van de partner bepaalt of hij geholpen wordt. Deze theorieën wekken de suggestie dat samenwerking gestabiliseerd wordt als mensen hun beslissing om een ander te helpen baseren op gedrag uit het verleden.

In **hoofdstuk 6** stel ik de vraag hoe directe en indirecte wederkerigheid tezamen menselijke samenwerking kunnen ondersteunen. Deelnemers aan het experiment konden er herhaaldelijk voor kiezen een ander te helpen ('helpen' betekent hier: een bedrag aan de partner overmaken), of om dat niet te doen. Ze konden deze beslissingen baseren op het gedrag van hun partner richting henzelf (directe ervaring uit interacties met een partner, waarop directe wederkerigheid kan worden gebaseerd), of richting andere deelnemers (reputatie-informatie, waarop indirecte

wederkerigheid kan worden gebaseerd). De resultaten wijzen erop dat, gemiddeld genomen, directe interacties zwaarder wegen in een besluit om te helpen. Interessant genoeg kunnen directe ervaring en reputaties elkaar compenseren in het in stand houden van samenwerking: als een partner voorheen geweigerd heeft te helpen in directe interacties, kan een goede reputatie toch nog tot samenwerking leiden. Andersom heeft een slechte reputatie relatief weinig effect, als de directe interacties maar positief waren. Dit wijst erop dat deze twee mechanismen niet onafhankelijk zijn; om hun rol in menselijke samenwerking goed te kunnen begrijpen, moeten onderzoekers zich richten op het samenspel van directe en indirecte wederkerigheid, en ze niet langer apart van elkaar onderzoeken. Een interessante bevinding is verder dat, net als bij de sociale leerstrategieën, mensen sterk van elkaar verschillen. Sommigen richten zich vooral op directe interacties, terwijl anderen hun beslissingen alleen baseren op de reputatie van hun interactiepartners.

In **hoofdstuk 7** ga ik dieper in op het idee van indirecte wederkerigheid. In hoofdstuk 6 werd dit begrip namelijk op een vrij simpele manier benaderd, simpeler dan hoe dit in de realiteit kan werken. Er wordt namelijk alleen maar naar het gedrag van een interactiepartner in het verleden gekeken, en niet naar de motivaties achter dat gedrag. Echter, hetzelfde gedrag kan totaal verschillende oorzaken hebben; een free-rider kan weigeren anderen te helpen omdat hij intrinsiek niet geneigd is tot samenwerking, terwijl een coöperatief individu een keer kan weigeren om hulp te geven aan een free-rider om hem daarmee te straffen. In hoofdstuk 7 onderscheid ik deze twee gedragsstrategieën van elkaar in een laatste experiment. Deelnemers kunnen nu ook meewegen of hun partners voorheen hulp gaven aan individuen met een goede of een slechte reputatie. Theoretische modellen laten zien dat het maken van dit onderscheid een sterk stabiliserend effect kan hebben op samenwerking in groepen; namelijk, als het bestraffen van free-riders wordt beloond, kunnen coöperatieve individuen onderling blijven samenwerken, en een goede reputatie behouden. De resultaten van dit experiment laten zien dat mensen sterk verschillen in de mate waarin zij belang hechten aan de motivaties van hun interactiepartners. Sommigen vinden het belangrijk dat een partner 'rechtvaardig' handelt, en bestraffen free-riders door ze nooit te helpen. Echter, lang niet alle mensen wegen de motivaties achter

het gedrag van hun partners mee in hun beslissingen: velen baseren hun samenwerkingsgedrag alleen maar op het gedrag van hun partner, en niet in de motivaties erachter; ze zijn niet geïnteresseerd of het gedrag van hun partner bijvoorbeeld was bedoeld om iemand met een goede reputatie te helpen, of iemand met een slechte reputatie te straffen. Het type individuen waaruit een groep is samengesteld kan een sterke invloed hebben op de dynamiek van samenwerking. Zo kan het bijvoorbeeld essentieel zijn om een aantal 'rechtvaardige' individuen in een groep te hebben om free-riders onderdrukken, en daarmee stabiele samenwerking in stand te houden.

Samenvattend laten de hoofdstukken 6 en 7 zien dat mensen sterk van elkaar verschillen in hun gedragsstrategieën. Sommigen laten hun samenwerking afhangen van directe interacties, terwijl anderen naar de reputatie van hun partners kijken. Grote diversiteit is ook terug te zien in de mate waarin belang wordt gehecht aan de motivaties achter het gedrag van hun partners: waar sommigen belang hechten aan de rechtvaardigheid van hun partners, speelt dit voor anderen geen rol in het aangaan van samenwerking.

Conclusie

De theoretische resultaten laten zien dat de culturele evolutie van samenwerking afhangt van de sociale leerstrategieën die mensen gebruiken. De ideeën uit de literatuur, waarin wordt betoogd dat culturele groepsselectie gemakkelijk tot de evolutie van coöperatie leidt, zijn voorbarig: de uitkomst van culturele evolutie wordt bepaald door het samenspel van sociale leerstrategieën en specifieke manier waarop groepsselectie plaatsvindt.

De experimentele resultaten geven inzicht in zowel sociale leerstrategieën als gedragsstrategieën in de context van menselijke samenwerking. Het opmerkelijkste aspect dat uit de experimenten naar voren komt zijn de grote individuele verschillen tussen mensen. Er bestaat grote variatie in zowel de manier waarop mensen van anderen leren, als in hun gedragsstrategieën in een context van samenwerking. Deze verschillen

kunnen grote consequenties hebben voor de uitkomst van interacties, en nader onderzoek is vereist om deze in kaart te brengen.

Sociaal leren heeft een grote invloed op menselijk gedrag. Om deze invloed te kunnen doorgronden, is een solide begrip van culturele evolutie vereist. Hiervoor is zowel theorievorming als gedegen empirische kennis van sociale leerstrategieën noodzakelijk. Uiteindelijk zal onderzoek naar het samenspel tussen sociaal leren en samenwerking leiden tot een dieper begrip van datgene wat ons mensen zo uniek maakt: ons sociale karakter.

Acknowledgements

Completing this dissertation would have been impossible without the help of many other people. First of all, I would like to thank my supervisor Franjo Weissing for his guidance throughout the process that resulted in this thesis. I very much appreciated the freedom he gave me to develop my ideas beyond the domain of theoretical biology, and his support in setting up the experiments on social learning, which excited me the most as a PhD student. His critical view and ever constructive feedback greatly stimulated me to keep developing myself. I am also thankful for the extensive amounts of time he took to comment on my writings, to reflect on my presentation skills and to give me advice on numerous aspects of scientific practice. All of these things have been very helpful to me in the past four years, and I am sure they will remain of great value in the future. Thank you Franjo, thank you for all the opportunities you gave me, it was a great pleasure to be your student!

My time in Groningen would have been a lot less interesting if it wasn't for Piet van den Berg and Andrés Quiñones. These guys made great contributions to the work presented in this thesis, and I learned a lot from both of them while working together on different projects. Dudes, I greatly enjoyed the honour of sharing an office with you, and I will really miss our daily exchanges about science, cultural evolution, football, quirky stress toys and Colombian pets.

The Theoretical Biology group was a great place to do my PhD, and many people in this group gave me very helpful advice in writing my thesis. I thank Ido for helping me out on numerous mathematical and statistical issues I could not get my head around; professor Pen, your sharp judgment and remarks considerably improved my work. I also benefitted from discussing various aspects of my research with Sander, Jordi and Alex; guys, I am grateful for your positive input. All other TheoBio's contributed in one way or the other to my work, not in the least by making my time in Groningen such a nice experience. Thank you Bram, Ana, Aniek, Ruth, Daan, Ivan, Leif, Andri, Sebastian, Jianlei, Chunyan, Joke, Kuke, Stefany and Ingeborg and for the good times!

Acknowledgements

Special thanks go out to Martijn Egas. He introduced me to research on human cooperation and how aspects relating to its evolution can be studied with decision making experiments. He instigated my enthusiasm for research and motivated me to pursue a PhD on this topic. I learned a lot from our ongoing collaboration, the results of which constitute two out of the six research chapters of this thesis. Martijn, thank you for all the work you put into our projects on reciprocity, and the time you took to discuss the designs, analyses and manuscripts related to them! I would also like to thank Eva van den Broek for supervising me as a Master's student at CREED in Amsterdam back in 2007, guiding me on the first steps towards setting up my own experiments, and showing me how much fun it can be to do research. I also thank Violet Swakman for her great work in one of the reciprocity projects; it was a pleasure to work with you! I also thank Aljaž Ule for encouraging discussions. I thank Mikael and Jaakko for the stimulating collaboration in projects that are still ongoing. Herrat, kiitos!


Here, I would express my gratitude to the members of the reading committee, Ton Groothuis, Dirk Engelmann and Manfred Milinski, who took the time to read critically through my thesis and give valuable comments.

My life outside of university was particularly fun due to some nice people I got to know during my time in Groningen. Of course the soon-to-be world famous Priceless Camels, Bram, Ivo, Maarten, Bob, Bert and Rangco, I had a blast rocking Leek and Appingedam, thanks for all! Vanesa, thank you for the good times, aberrant conversations and broken coasters. Thijs, I am grateful for the loads of programming advice, but even more for the fun times drinking beer and playing chess. I also would like to thank some older friends who were essential for me keeping a positive spirit during my PhD: Tom, Ferdi, Chantal, Veerle, Freek, Koen, Bas, Rik, Bart, Daniël, Karst, Albert, Eva, Jarno, Ruud and Koen, thanks guys!

I would like to thank Thijs (yes, again) and Toon for their willingness to be my paranymphs. I look forward to see you two in your suits. I thank my parents, and my three brothers for their care and support along the way. Last but certainly not least I thank Suzanne for being my all-round advisor, but even more for always being there for me! Thank you!!

Biography

Lucas Molleman was born in 1984 in the Dutch village of Zieuwent. After the gymnasium in Groenlo, he studied Biology and Philosophy at the University of Amsterdam. In 2009 he obtained an Ubbo Emmius grant to pursue his PhD in the Theoretical Biology Group at the University of Groningen, The Netherlands. His research focuses on decision making, human cooperation, cultural evolution and social learning. In November 2013 he continued his work as a postdoctoral researcher at the Centre for Decision Research and Experimental Economics at the University of Nottingham in the United Kingdom.



(photo by Suzanne Hoogers)