On the reliability of unreliable information: Gossip as cultural memory

Dominic Mitchell[1,*], Joanna J. Bryson[1,2,*], Paul Rauwolf[1,3], and Gordon P.D. Ingram[4]

[1]*Department of Computer Science, University of Bath, UK*

[2]*Center for Information Technology Policy, Princeton University, USA*

[3]*Institute for New Economic Thinking, University of Oxford, UK*

[4]*Department of Psychology, Universidad de Los Andes, Bogotá, Colombia*

June 6, 2016

Abstract

When individuals learn from what others tell them, the information is subject to transmission error that does not arise in learning from direct experience. Yet evidence shows that humans consistently prefer this apparently more unreliable source of information. We examine the effect this preference has in cases where the information concerns a judgment on others' behaviour and is used to establish cooperation in a society. We present a spatial model confirming that cooperation can be sustained by gossip containing a high degree of uncertainty. Accuracy alone does not predict the value of information in evolutionary terms; *relevance*, the impact of information on behavioural outcomes, must also be considered. We then show that once relevance is incorporated as a criterion, second-hand information can no longer be discounted on the basis of its poor fidelity alone. Finally we show that the relative importance of accuracy and relevance depends on factors of life history and demography.

Keywords: gossip, reputation, cooperation, social norm

*d.mitchell@bath.ac.uk, jjb@alum.mit.edu, *Corresponding author; p.rauwolf@bath.ac.uk, gp.ingram@uniandes.edu.co

1 Introduction

A cooperative society is defined as one in which individuals benefit from the collective absence of defection (Axelrod and Hamilton, 1981). For any individual member defection is advantageous only if it is the exception. An example of cooperation is free trade between nations (Schelling, 1960). Against a context of free trade any one nation will benefit by introducing import tariffs but if all do so, all lose (Hardin, 1968). This paradox makes cooperation vulnerable, thus the discovery of cooperation in many different animal societies is surprising. Evolution depends on the relative fitness of individuals so we might expect individuals to maximize this regardless of its consequences for the group (Williams, 1966; Dawkins, 1976). The existence of cooperation under such conditions can be explained by direct reciprocity (Trivers, 1971): if interactions in the society are between those who have met before and will meet again, a you-scratch-my-back rule between pairs results in a society which meets the criteria for cooperation. However cooperation also exists in societies such as contemporary urban ones, where populations are large and mobile and interactions are between strangers that are unlikely to meet again. Here direct reciprocity fails as an explanation, as the society is vulnerable to free-riders—individuals who defect opportunistically and move on before the consequences of their behaviour can catch up with them (Enquist and Leimar, 1993). Although such lifestyles of course have their own costs and benefits (Taylor, 2014), in such cases a different mechanism is typically used to explain cooperation. *Indirect reciprocity* (Alexander, 1987; Nowak and Sigmund, 1998, 2005) is a system in which agents do good to one other because they are likely to receive good from a different other. One way this can be achieved is when individuals observe each other's behaviour and pass on the resulting reputation in social transmission—gossip. Defectors can no longer free ride; however mobile they are they risk being preceded in their interactions by their reputation (Milinski et al., 2006).

In relationships based on direct reciprocation we can imagine that it is quite possible for an individual to form a judgment of the quality of a partner because the history of previous interactions provides its own guarantee. But in order for a system of indirect reciprocity to hold, a similar quality of information is required of partners who have never met before (McElreath and Boyd, 2007). Requiring a reputation to achieve indirect reciprocity places greater demands on it (Roberts, 2008). Whereas direct reciprocity requires memory of previous experience of direct perception, here the information is a temporally and spatially displaced judgment determined and passed on by others and as such potentially more vulnerable to error. In this article, we focus on the question of how gossip can support cooperation (Nowak and Sigmund, 2005) despite this potential for inaccuracy.

In cases where the reputation of a partner is learnt third-hand via social transmission of information (gossip) the reliability of that information is paramount. We might expect individuals to be very sensitive to any evidence of error in the information. However a number of researchers have shown that our behaviour contradicts this expectation. Human societies continue to be trusting of socially learned information even when it is known to contain much error and uncertainty (Richerson and Boyd, 2005; Hess and Hagen, 2006). What is more, when faced with the direct evidence of our senses contradicting social learning we often give preference to the latter (Asch, 1956) even when based on anonymous gossip (Sommerfeld et al., 2007). Comparisons of human and non-human cognitive abilities (Nagell et al., 1993; Whiten et al., 2009) cite humans' greater use of culture—which is founded on social learning—as our key advantage (Gergely and Csibra, 2006; Henrich and McElreath, 2003; Tomasello, 1999). Human infants when taught to perform a task in an inefficient manner persist in it despite exposure to a better alternative (Gergely et al., 2002; Meltzoff, 1988), while other primates imitate only necessary actions once some demonstrated actions are recognised as unnecessary (Whiten et al., 2009). In this article we present a possible explanation for the human propensity to over-imitate, by analysing the information in gossip relative to information from direct experience.

Humans intuitively value accuracy in information and have a variety of methods to assess it. We assign lower value to information that is contradictory, or that differs from an authoritative source. Scott-Phillips (2008) argues that in nature information should be defined in evolutionary terms. But evolution does not access the accuracy of information directly, it acts on the results of the information being put to use (Rauwolf et al., 2015). This has the effect of introducing *relevance* (Dessalles, 2007; Sperber and Wilson, 1995) as a crucial component of information value. We define relevant information as information that is needed to perform the task at hand. Information that is highly accurate but unlikely to be put to use will therefore have less value. Non-relevant information will not be acted upon, and therefore will not be subject to selection pressures, whether it is accurate or inaccurate.

When individuals rely on experience, we assume accuracy is not in question but the relevance of what they learn is dictated by the limits of the life history of the individual concerned (Čače and Bryson, 2007; Roberts, 2008). In the case we examine—the role of reputational information in successive interactions—the information in question can only be put to use when an individual meets the same individual again and the consequences of reputation gained earlier are realized. The notion of relevance in this context is captured by the cost / benefit relationship in the following equation:

$$c < wb \tag{1}$$

where w is the probability of future interactions, c is the cost incurred at the first interaction, and b is the benefit expected from that cost in any future interactions (Roberts, 2008). When an individual learns by social communication, relevance is likely to be higher as social transmission decouples the information from full dependence on physical experience. In the limit case of instant communication, $w = 1$, agents can be sure to know the most recent social actions of any individual they meet. This does not mean that social information is always of relevance to the audience of gossip, but that on the occasions when an interaction is influenced by the reputation of an agent there is a higher probability of benefit from information if it is socially transmitted than if it must be learnt through individual experience. This is because it is more likely that others may have had recent (or indeed any) experience of the interactor. However in freeing information about interactions from spatial and temporal constraints, socially transmitted information may incur a cost due to loss of accuracy. We are agnostic here as to whether error is introduced deliberately or by misperception; we only assume that error accumulates over the process of repeated transmission. We can again describe a necessary relationship:

$$q < wb/c \tag{2}$$

where q is the probability of the reputational information being correct. That is, the expected future benefit must outweigh the cost of the interaction by a sufficient extent to compensate for the probability of miscommunication.

In a series of three experiments we explore the effect of incorporating both relevance and accuracy in the evaluation of information learned by direct experience or by social communication. First we replicate in a spatial simulation the mathematic result of Nakamura and Masuda (2011). This demonstrates that in the full social communication case ($w = 1$) reputational information can sustain cooperation despite containing a large degree of uncertainty caused by assigning others reputations based on incomplete observation. In Experiment 2 we relax the ($w = 1$) assumption by parameterizing the rate of spread of social information, and substituting low transmission fidelity for incomplete observation as the cause of error. We show that accuracy alone cannot predict the ability of social communication to sustain cooperation. A level of accuracy too low to sustain cooperation at a lower rate of transmission is nevertheless able to do so if the transmission rate is increased. In the third experiment we replace control of transmission rate with the assumption of a limited lifespan for agents. We compare the performance of individuals that learn from direct experience with those that use social communication. A limited lifespan means that reputational information is liable to be outdated and thus of lower relevance. But as the direct experience strategy can be more vulnerable in this respect we again show that accuracy alone cannot predict the outcome. A gossiping strategy which is ruled out by its high level of error under conditions of long or unlimited lifespan may still outperform experience-based learning if lifespans are shorter. Our conclusion is that the relative merits of relevance and accuracy are highly sensitive to the ecology of the society. Adhering to a gossiping strategy despite the potential inaccuracy of the information can thus be evolutionarily successful under many conditions.

2 Methods

Model

All experiments are based on the same agent-based model of evolutionary game theory. Interaction between individuals following different strategies and influenced by different kinds of information generates varying payoffs which are used to evolve the population.

Individual Interactions: the Donation Game

Following the method used by Nakamura and Masuda (2011), agents in our simulation play the donation game, which is similar in form to the prisoner's dilemma (See Figure 1). The idea of the game is based on ordinary social interactions, where an individual can choose to help a passing stranger or not, for example by giving directions. Individuals move at random in a two-dimensional world at each iteration. The world is a 21 x 21 torus which wraps horizontally and vertically. If two agents are located adjacently, roles for the game are assigned randomly and they are able to interact by playing a move in the game. The game has two players: a donor and a recipient. There are two possible moves in the game. If the donor cooperates, it pays cost c and the recipient gains benefit b. If the donor defects it pays nothing and the recipient gains nothing. In cases where $0 < c < b$ the dilemma of cooperation described in the introduction is created: in a single interaction the best move for the

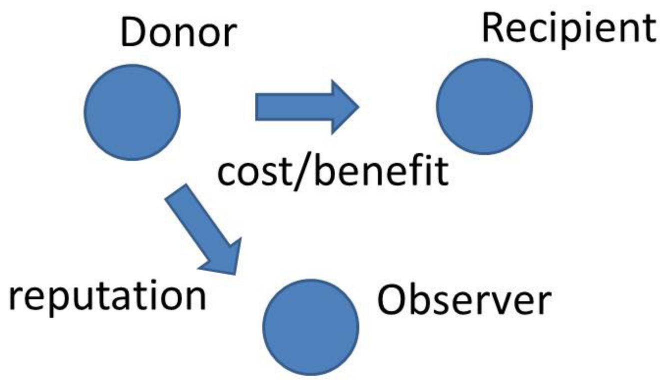

Figure 1: The donation game involves three actors: a *donor* who chooses whether or not to donate to a *recipient*, and an *observer*. Where a donation occurs, the donor pays c and the recipient receives b. The observer updates its own record of the donor's reputation both on the basis of the donor's action, and possibly on the reputation the observer holds for the recipient (see main text and Table 1.)

Table 1: Social norm used in experiments 2 and 3. The second-order judging norm.

	Good	Bad
Cooperate	G	B
Defect	B	G

donor is to defect, but each individual does better if all cooperate.

Reputation Assigned by Observers According to a Social Norm

Individuals have binary reputation values: Good or Bad. The observer updates the donor's reputation based on a combination of the donor's action: Cooperate or Defect, and the recipient's reputation according to a single model-wide rule called the social norm. Here we use the judging norm (Nakamura and Masuda, 2011; Ohtsuki and Iwasa, 2006, see Table 1). This specifies a society which approves of those who cooperate (C) with the good (G) and defect (D) against the bad (B). Each individual has its own record of others' reputations. Individuals meet at random and update each other's records so that reputation spreads. For a discussion of other possible norms and their effects on donation games, see Rauwolf et al. (2015).

Reputation Impacts on Individual Choice of Action

The donors' choice of action is determined by their action rule (Nakamura and Masuda, 2011). There are 3 populations each acting according to a different rule: unconditional cooperators (ALLC agents) cooperate regardless of their parter's reputation; unconditional defectors (ALLD agents) defect regardless of their partner's reputation, and discriminators (DISC agents) cooperate with partners who are good but defect against bad partners (Ohtsuki et al., 2009).

Evolution According to Payoff

The population is reproduced according to fitness-proportionate selection (Goldberg and Deb, 1991) based on pay-offs from the donation game. Mutation, in which an action rule different to the parent's is substituted in offspring, occurs according to a probability (see Table 2).

Experiment 1. Judgments Made Without Full Knowledge

It can be seen from Table 1 that the observer's assignment of reputation to a donor is a second-order judgement. It is determined by a combination of the donor's action with the recipient's reputation. Nakamura and Masuda (2011) introduce uncertainty in the second term, so that a proportion of donors take part in and are judged on interactions with recipients of unknown reputation. The authors identify three social norms which are still able to realize indirect reciprocity under this condition. We replicate this in the case of a single norm—the suspicious judging norm—since rather than considering the relative performance of different norms we are interested in the extent to which a given norm can continue to function under conditions of unreliable information. We choose the suspicious norm motivated by the finding in Yoo (2009) that negative gossip has a greater impact on the listener.

Method

The donor and observer only know the reputation of the recipient with probability $0 < k < 1$. This is concomitant: all agents share the same information, so either they both know or neither do. In view of this the action rules are extended. ALLC agents here also cooperate (C) with partners whose actual reputation is unknown (U) to them. ALLD agents defect (D) against unknown partners. DISC agents cooperate with unknown (U) partners. The social norm is also extended (see Table 2). Players who cooperate with unknown partners are judged good (G) but those who defect against an unknown partner are judged bad (B). This norm is termed suspicious because faced with an agent who defects against a partner unknown to him observers have a choice of two assumptions. The trusting assumption is that he defects because the unknown agent is in fact bad. The suspicious one is that he defects despite the fact that his partner is in reality good.

Simulation

In each round, two randomly selected agents are assigned the roles donor and recipient and play the donation game. Payoffs to both players are determined according to the rules of the game described above. To begin by considering the most extreme case (this will be refined below), reputation is spread instantly and with perfect transmission fidelity to all agents resulting in a single global reputation for each agent which is then used to play the next round of donation game. An assumed observer updates the donor's reputation in the list held by all agents according to the social norm described above. The list records each agent as either good or bad[1]. Rounds are repeated until every agent has played once against every other. We call this one generation. At the end of each generation the sum of all the costs and benefits to each agent are calculated.

A generation is followed by a round of reproduction. The strategy of each agent is encoded in a genome of one position with three alleles. Each member of the new generation is selected from the previous generation according to fitness-proportionate selection. The chance of selecting individual i is $p_i/(\sum_{j=1}^{n} p_j)$, that is the payoff of individual i divided by the total pay-off of the population. There is also a small chance for mutation $u = 0.01$, where an agent's strategy is replaced by one of the three available strategies. At the start of the run the population is made up entirely of DISC agents. This process iterates 2000 times to provide ALLD and ALLC sufficient opportunity to invade DISC from scarcity.

[1]In experiments 2 and 3, global reputation is replaced by peer-to-peer gossip and individual memory in order provide a setting in which claims made about cheap communication can be examined.

Table 2: The suspicious social norm, which is used in Experiment 1. The judging norm is extended to include a first-order norm in cases where the recipient's reputation is unknown. The norm assumes that when a donor plays Defect against a partner with unknown reputation he is playing Defect to a Good partner, hence is suspicious of him, rather than the alternative possibility that he is playing Defect because his partner's reputation is Bad.

	Good	Bad	Unknown
Cooperate	G	B	G
Defect	B	G	B

Table 3: Table of free parameters, values used in present figures, and a sensitivity report including range of values tested.

Symbol	Parameter	Value	Notes
n	population size	100	number of agents.
r	rounds	10,000	donation games per generation.
g	generations	5,000	number of evolutionary iterations.
c	cost	1	cost of cooperating, kept constant.
b	benefit	[1,10]	benefit of being recipient of cooperating donor, increments of 0.25
u	mutation rate	0.01	chance of strategy mutation.
k	knowledge	[0,1]	probability recipient's reputation known, increments of 0.025

Results and Discussion

In the unshaded parameter region in Figure 2 interactions are entirely cooperative. A population which is predominantly DISC in the context of these parameter ranges and the suspicious judging social norm produces a society where interactions are overwhelmingly cooperative (C). The result is consistent with that of Nakamura and Masuda (2011). Cooperation is stable against invasion down to low values of k provided that the cost/benefit ratio is high. Conversely in regions where the cost/benefit ratio is low a higher value of k is required to enable cooperation.

These results suggest that within a specified context of action rules and social norms cooperation can persist in a society even when a significant amount of the information used to enforce the cooperation is uncertain. The assumption of trust specifies a cooperative action from the donor when the recipient's reputation is unknown (Nakamura and Masuda, 2011). This means that a certain proportion of recipients whose actual reputation is Bad will escape the punishment. That cooperators are able to overcome such limitations on observation demonstrates that it can be in an individual's interest both to tolerate being judged and to judge others, even in this imperfect fashion.

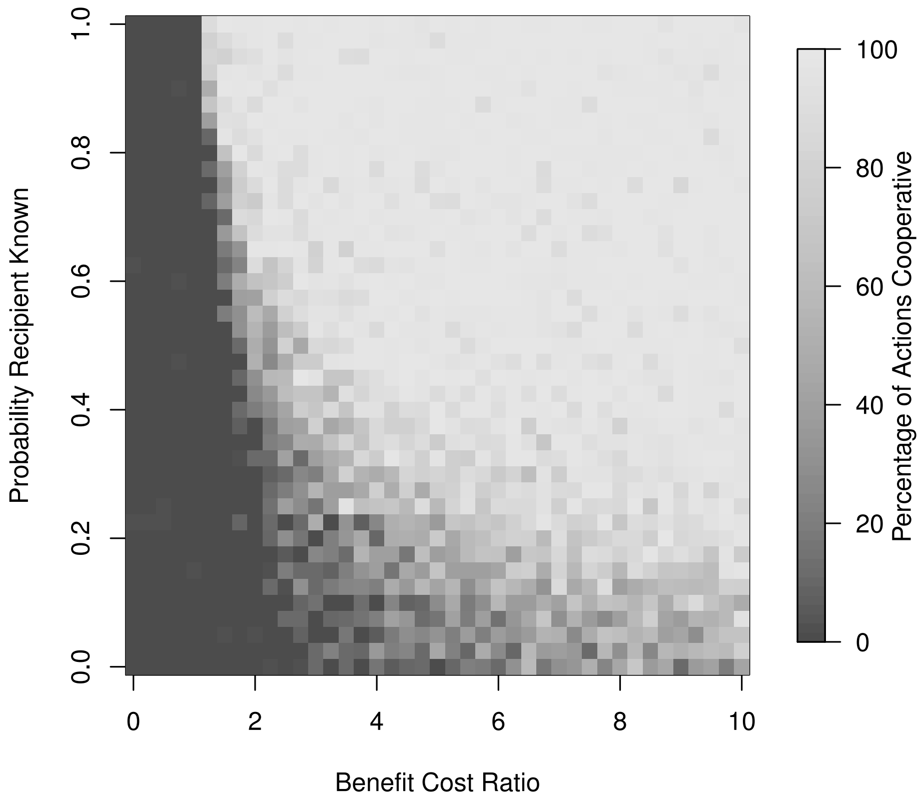

Figure 2: Judgments made without full knowledge. Percentage of cooperative interactions aggregated over 10 runs. In the unshaded parameter region interactions are cooperative. Dark cells indicate that at 2000 generations DISC is invaded by ALLD.
b

Experiment 2. Agents' Knowledge of Reputations Differs According to the Spread of Gossip

In Experiment 1 the information carried by gossip was sometimes incorrect at the source but there was no error in gossip transmission. But wee know that in the real world, knowledge and behaviour take time to diffuse through a population (Gruber et al., 2015). Here we investigate error in transmission and the speed at which information is spread. We replace the idealized broadcast transmission of Nakamura and Masuda (2011) with a more ecologically plausible spatial model of peer-to-peer interactions occurring during encounters between mobile agents.

Method

Agents interact when they are in sufficient proximity to each other. Specifically the world is an 21x21 grid, and agents interact on each cycle with at most one other agent, if one is present on an adjacent grid cell. After playing a round of the donation game the partners exchange information. This communication takes place regardless of the nature of the interaction that has just taken place. If one agent believes a third party to be bad, and the other believes it to be good, then the agent with the 'good' association updates this instead to 'bad'. We use this pessimistic rule to represent a society in which reputation is difficult to achieve but easy to lose.

In Experiment 2, the explicit probability used in Experiment 1 $(1 - k)$ that a recipient in an observed game may be unknown (U) is replaced by the emergent chance that information about a player's reputation may not have spread sufficiently fast to inform the donor or observer in their next interaction. All agents are initially assigned a good reputation, we examine how efficiently evidence of infringed reputation is spread. Therefore, agents no longer require individual behaviour rules for handling unknown (U) individuals. The probability of information being up-to-date is a function of the agents' collective behaviour. The rate at which gossip spreads is a function of a limit t placed on the number of individuals whose reputations may be exchanged at each encounter, see Table 2. This limit combined with the absence of a global reputation means that agents do not know each other's knowledge state and therefore cannot selectively divulge full information. Rather, agents communicate a random set of their own knowledge. Note that it is possible for the gossip to have reached the donor but not the observer or vice versa. This type of observation is known as *independent*, in contrast to Experiment 1 where due to the use of broadcast reputation observation was concomitant.

An error rate e specifies the probability that at each exchange of gossip the opposite update to the one intended is made by the recipient (see Table 2.) Populations of the three strategies (ALLC, ALLD, DISC) are of equal size throughout Experiment 2. There is no reproduction, selection or mutation. The relative performance of strategies is judged on their average pay-off from the donation game

Results and Discussion

Figure 3 shows that minimum values of both transmission rate and accuracy are required in order for DISC to out-perform ALLC and ALLD. However within this context there is an interplay between the

Table 4: Table of free parameters, values used in figure 3, including range of values tested.

Symbol	Parameter	Value	Notes
n	population size	150	number of agents.
r	rounds	1,000	spatial model so number of interactions per round varies
c	cost	1	cost of cooperating, kept constant.
b	benefit	3	benefit from being recipient of cooperating donor, kept constant
e	error	[0, 100]	per cent chance incorrect transmission of reputation
t	transmission	[0, 100]	percentage of population gossiped about at each encounter.

two parameters. A level of accuracy poor enough to cause DISC to fail at low transmission rates is nevertheless sufficient to make the DISC the winning strategy once transmission rate is raised.

When accuracy and relevance are both factors in the value of information, socially-learned information even with a high incidence of error may outperform experience-based learning if the latter is insufficiently relevant (Sommerfeld et al., 2007). By parameterizing the rate at which information spreads, we are able to examine the interaction between relevance and accuracy. Although relevance is not directly controlled in our agent-based model, as transmission rate increases so does the probability of having reputational information concerning an individual that an agent is interacting with. The risk of being updated by inaccurate information is outweighed by the risk that not being updated will cause greater inaccuracy.

Formal analysis of an ALLD agent's reputation

To better understand how transmission and error rate affect the performance of each strategy, we turn to formal analysis. Cooperation is maintained by eliminating free-riders, and this is only accomplished if free-riders are negatively reputed. Here we analyse the probability, $ALLD_b$, that an ALLD is considered bad.

To do this formally, we presume a few simplifications. Gossip and the spread of social information inherently rely on the specifics of a social network. If an individual sees one person more than another, they are more likely to receive information from that person. While computer simulations easily permit spatial and network dependencies, here we focus on elucidating the interplay between transmission and error rate. In doing this, the nuances of a particular social network are abstracted away. However, as a consequence we slightly alter the definition of error rate and transmission rate.

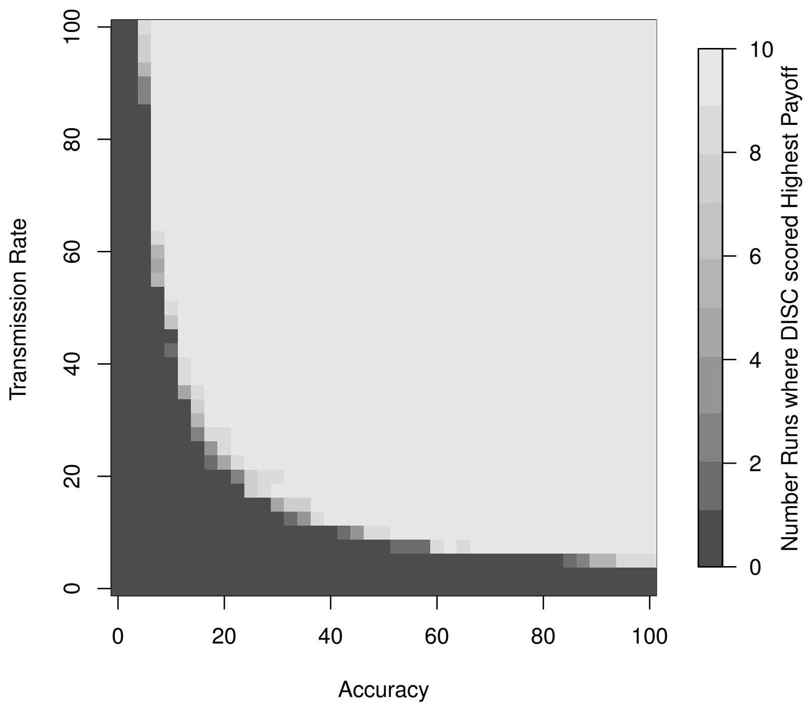

Figure 3: Experiment 2, in which agents' knowledge of reputation differs as a result of the spread of knowledge. Simulation results are shown for a range of error rates in transmission (e), and rates of transmission (t), aggregated over 10 runs. In the light region DISC scores a higher payoff than both ALLC and ALLD.

When an agent uncovers a free-rider, it attempts to pass on this information through gossip. We define transmission rate, t', as the probability an agent will hear the newly discovered information during a round of play. Previously, transmission rate was defined as the chance that two interacting agents would gossip about another agent. The gossip could then be passed to other agents at some later interaction. However, the fact that these simulations are spatial—that is, each agent has a location and this determines which agents they interact with—has consequences. While transmission of gossip relies on whether or not a bit of information is discussed, it also relies on the probability of interacting with other agents with relevant knowledge. This is correlated with location. Here we define transmission rate to encapsulate both space and the chance to spread gossip. t' is the chance that, in any round, an agent will hear a specific piece of gossip.

When an agent hears gossip there is a chance it will be in error. e', is the chance communication is erroneous. Again, this is slightly different compared to the error rate employed in the computer

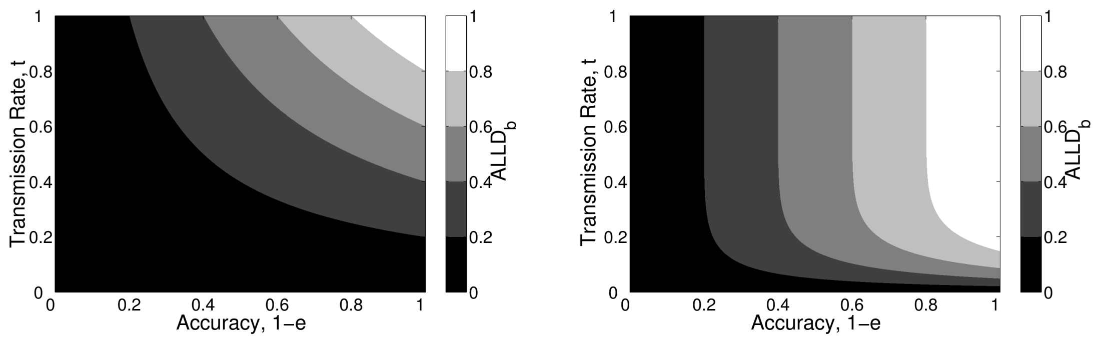

(a) one round (b) ten rounds

Figure 4: The probability $ALLD_b$ that an agent has heard that a particular ALLD agent is bad, after **(a)** one and **(b)** ten rounds of play. Mathematical outcomes are shown for a range of error rates in transmission (e') and rates of transmission (t').

simulations. Previously, error rate represented the chance that at a given communication was incorrect. Here we presume that whenever an agent hears information, there is an $1 - e'$ chance that it is correct.

If gossip travelled instantaneously, then $1 - e'$ fraction of agents would correctly diagnose the defector. However, information does not spread instantaneously, it depends on the speed of transmission. So, if an agent discovers a defector, the chance this information will accurately reach another agent in a given round is:

$$ALLD_b = (1 - e')t' \tag{3}$$

An ALLD is marked as bad if the agent receives the information, t', and the information is correct, $1 - e'$. Furthermore, the probability that an agent will accurately hear the gossip after i rounds can be written as:

$$ALLD_b = (1 - e')(1 - (1 - t')^i) \tag{4}$$

There is a $1 - t'$ chance that an agent does not hear the information in a given round. Consequently, there is a $(1 - t')^i$ chance that an agent has not heard the information after i rounds. $1 - (1 - t')^i$ is then the probability that an agent has heard the information in i rounds. This result is then multiplied by the probability that the information is accurate, $1 - e'$.

Results and Discussion

A parameter sweep of e' $[0,1]$ and t' $[0,1]$ can be plugged into Equation 4 to calculate the likelihood than an ALLD agent is given a bad reputation, $ALLD_b$. Figure 4(a) shows $ALLD_b$ after one round of gossip. Clearly, $ALLD_b$ depends on the interplay between the transmission and error rate. As the transmission rate increases, so does the likelihood of speedily informing others. However, the accuracy of the transmission is limited by the error in communication.

Figure 4(b) shows the percentage of negatively reputed ALLD agents after 10 rounds of gossip. As time passes, the effects of the transmission rate diminish. After ten rounds, an agent inhabiting a relatively low-error-rate world is likely to have heard the information, regardless the transmission rate. Only very small levels of t' show a marked difference.

However, whilst the effect of the transmission rate may become moot after several rounds, cooperation might not last that long. If an ALLD is permitted to operate with a good reputation for multiple rounds, then it might undermine a cooperative society long before future rounds can be played. This is why, in Figure 2, the strategy with the highest pay-off depends on the rate of transmission. Expedited and relevant information dissemination reduces the opportunity for free-riding, increasing the pay-offs of DISC agents.

Experiment 3. Personal Experience Versus Gossip Under Limited Lifespan

In Experiment 2 the value of information about reputations did not degrade over time. The information an individual has may be incorrect either because of transmission error or because restriction on the transmission rate prevents incorrect information from being updated. But there is another way in which information could be false which was not considered: that is, in the time since the information originated the object of its reference may have changed. Here we revoke this assumption and give the value of the reputational information a time limit. Information that was originally relevant may subsequently become irrelevant when the state of affairs it applies to has changed. We do this by imposing a limited lifespan on the agents, affecting both the individuals that carry information and the individuals that the information refers to.

Table 5: Table of free parameters, values used in Figures 5 and 6, and values tested.

Symbol	Parameter	Value	Notes
n	population size	200	number of agents.
r	rounds	5,000	spatial model so number of interactions per round varies
e	error	[0.7]	per cent chance incorrect transmission of reputation
l	lifespan	[3000, 10000]	at initiation each individual's life set at random random between 0 and this upper limit

Method

In this simulation we compare two strategies: *Direct Observers* and *Gossipers*. Direct Observers do not transmit gossip and do not receive it. Their only source of information is their own observations. Here we assume that observer's information on others' reputations is perfectly accurate, being based on their own (perfect) perception, and not subject to transmission error e. This assumption is the limit case, giving the observers the maximum advantage of their strategy. In contrast, Gossipers transmit and receive gossip about others' reputations. The accuracy of their information is affected by the value of e and the number of gossiping interactions that the information has crossed.

We introduce a lifespan value l that determines the maximum number of rounds an agent lives. Whether lifespan is a genetically regulated trait is a matter of some controversy, lifespan is a common feature used to characterise species, and from a theoretical perspective it has known impacts (Kokko, 1997; Čače and Bryson, 2007). At the beginning of a simulation, each agent is randomly assigned an age within the lifespan. After each round the lifespan reduces by 1. When an agent reaches the end of its life it is deleted. Dead agents are replaced by new agents which possess no information from either gossip or observation. Other agents in turn have no information from direct observation or gossip regarding the reputation of new agents.

As in Experiment 2 at the start of the run agents' reputations are set to good (G) and gossip consists of setting good (G) to bad (B) only. This provides a simple rule for which agent takes precedence in information exchanges. As in Experiment 2, populations of all six strategies are of equal size. Note that as in Experiment 2 there is no selection or mutation, even though there *is* reproduction. New agents are identical clones of their parent to limit the experiment to testing the effect of irrelevance generated by agent turnover. The relative performance of Direct Observation versus Gossip is judged on the extent to which each population has assigned bad reputations (B) to ALLD (defector) agents.

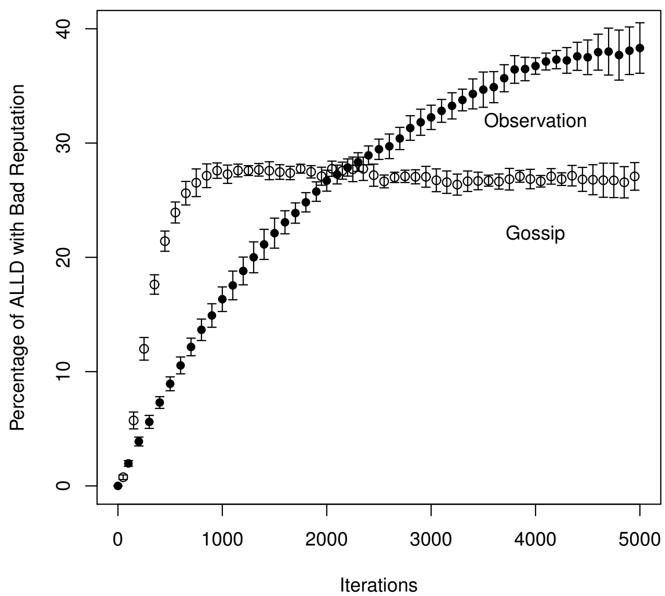

Figure 5: Direct Observation versus Gossip under a longer lifespan, $l = 10000$, $e = 0.7$, the presence of an error rate in gossip means that over time it can be outperformed by direct observation. Points plotted are means with 95% confidence interval calculated over ten replications.

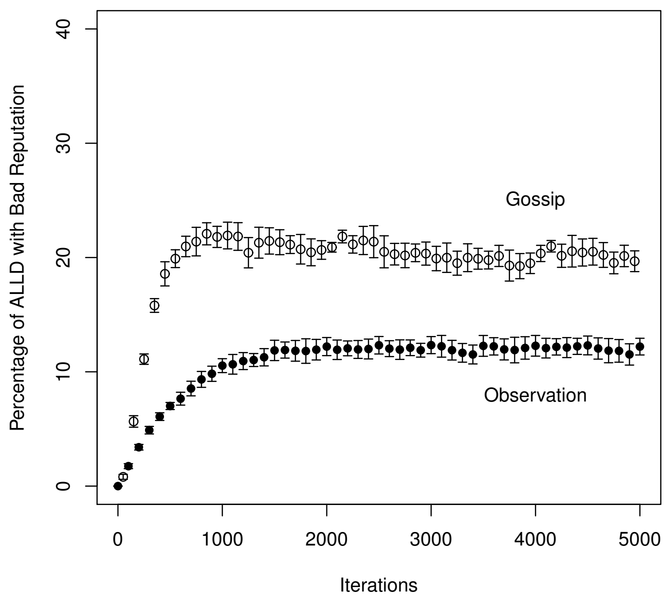

Figure 6: Direct observation versus gossip under a shorter lifespan $l = 3000$). Despite a ceiling determined by the error rate $e = 0.7$ here gossip can perform better than direct observation. This is because knowledge from observation will on average be acquired on a lower number of individuals per turn, therefore it suffers higher impact from the lower number of opportunities per lifetime. Thus the error generated during social transmission does not necessarily entail that direct observation is a better strategy. Points plotted are means with 95% confidence interval calculated over ten replications.

Results and Discussion

From Figure 5 it can be seen that, for the conditions of e and t in our experiments, gossiping individuals can initially learn more quickly than those that rely on personal experience. This is in keeping with the results demonstrated in Experiment 2. However transmission error acts as a ceiling on the maximum amount of knowledge acquired from gossip—after initial progress a stage is reached where updates are no longer improving on the information that an individual already holds. Learning by direct observation may be slower, but since it is not subject to transmission error it can eventually deliver a greater amount of information to each individual. However, this effect depends on lifespan which limits what one agent can discover. Figure 6 shows populations competing under the same conditions except for a shorter lifespan. Both strategies show a reduction in performance as a more changeable ecology places greater demands on information-gathering abilities and reduces the relevance of information held. However the two strategies do not degrade equally—for this set of parameters, the advantage of accuracy over relevance has been overturned. Gossiping individuals have access to more information as the error caused in transmission is outweighed by the speed advantage of gossip.

When lifespan is limited, relevant reputational information levels are reduced in four ways:

1. a record of reputation is useless when the individual it refers to no longer exists.
2. when individuals die the information they carry dies with them.
3. newborn individuals do not have a history of interactions on which their reputation can be based.
4. newborn individuals when entering into an interaction do not have the benefit of a history of personal experience or gossip to draw on when choosing which move to play.

If information spreads more quickly by gossip than it can by direct observation, then gossip meets these challenges more effectively, providing only that the rate of error is sufficiently low compared to the relative rate of transmission. We formalise this provision below, and show results over a larger range of parameters.

Formal Analysis of Gossip versus Direct Observation

The utility of gossip and experience can be formally analysed. By comparing the ability of Direct Observers and Gossipers to negatively repute ALLD agents, we can analyse the contexts where Gossiping is a dominant strategy.

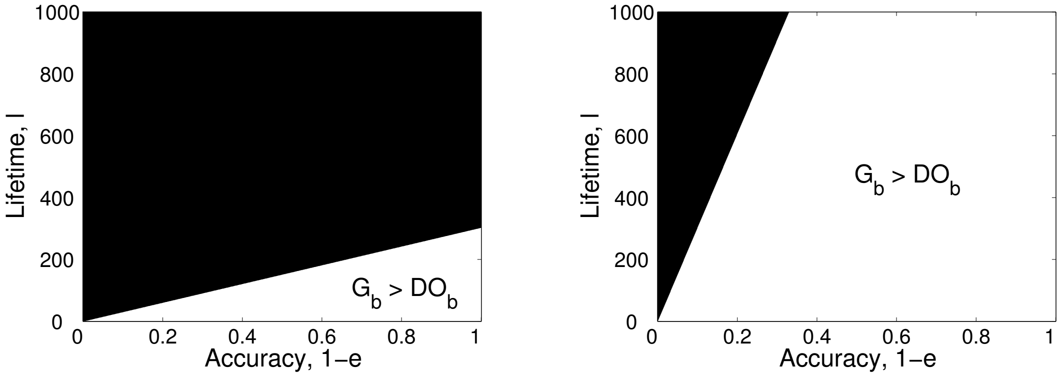

(a) $N_d = 50$

(b) $N_d = 500$

Figure 7: White represents the contexts where Gossipers diagnose more ALLD agents compared to Direct Observers. Analysis is run over a parameter sweep of lifespan (l) and error rate (e). Because lifespan does not scale with the size of the population, results vary with the latter. The number of ALLD agents (N_d, held constant at $1/3N$) in figures shown is **(a)** $N_d = 50; (N = 150)$, and **(b)** $N_d = 500; (N = 1500)$.

What is the probability that a Direct Observer will negatively repute an ALLD agent? As before, each Direct Observer begins with the belief that all agents are good. It is only through experience that agents attribute bad reputation. We define l as the Direct Observer's life expectancy. l denotes the number of agents it will meet during the course of its life. For simplicity and without loss of generality, we presume that at every interaction the agent meets a previously unmet agent, thus again maximising the benefit of the Direct Observation strategy. Therefore at the end of its life, the agent will have met l agents — one for each interaction.

During an interaction, the agent will meet an ALLD with probability P_d, where P_d is the percentage of ALLD agents in the population. Since the agents are judging others by direct experience, we assume no error in the observation. By the end of its life these optimal Direct Observers will have correctly classified lP_d ALLD agents as bad. The fraction of negatively reputed ALLD agents is this total divided by the number of ALLDs in the population. We define N_d as the number of ALLD agents in the population.

We presume that at any given time the age of the agents are homogeneously spread between the range $[1...l]$. Thus, the average age of the population is $\frac{l}{2}$. If DO_b is defined as the average fraction of negatively reputed ALLDs classified by Direct Observers, then it can be written as:

$$DO_b = \frac{lP_d}{2N_d} \tag{5}$$

In contrast to Direct Observers, Gossipers help each other classify ALLDs. As a consequence, while still constrained by l, the Gossipers benefit from a multiplier t which is the number of individuals they learn about on each cycle. However, they must contend with the possibility of erroneous communication.

Again for the benefit of simplicity, we assume a well-mixed population, such that in the first round, each ALLD is observed by at least one Gossiper. This also assumes that there are more Gossipers than ALLD, a safe assumption if half of the population is a Gossiper and a third ALLD. The Gossipers then relay the information to others. The likelihood that a Gossiper negatively reputes an ALLD is calculated as per Equation 4. The only undefined parameter from the equation is the number of rounds the agent has experienced, i. If each agent can experience l interactions in its lifetime, then, on average, a Gossiper will have experienced $\frac{l}{2}$ interactions. $i = \frac{l}{2}$ is then the average number of rounds the Gossiper has experienced. If G_b defines the chance that a Gossiper negatively reputes an ALLD, then:

$$G_b = (1 - e')(1 - (1 - t')^{\frac{l}{2}}) \tag{6}$$

(Note that $(1 - t')$ is raised to $\frac{l}{2}$, one half of lifespan.)

Now we can calculate when Gossipers will diagnose more ALLD agents than Direct Observers. Namely, when is:

$$G_b > DO_b \tag{7}$$

a somewhat complex relationship that depends on the interactions between e', t' and l

$$e' < 1 - \frac{lP_d}{2N_d * (1 - (1 + t')^{l/2})}. \tag{8}$$

Results and Discussion

Figure 7 shows when Gossipers outperform Direct Observers for differing values of lifespan, $l = [1, 1000]$, and error rate, $e' = [0, 1]$. As in the computer model, 1/3 of the population is comprised of ALLD agents (i.e. $P_d = 1/3$). The transmission rate is held static at $t' = 0.5$. As explained in Experiment 2, because $i = l/2$, i is typically large enough that the transmission rate does not affect the fraction of negatively reputed ALLDs.

Because Direct Observation relies on personal experience over a fixed lifespan, the winning strategy is not only dependant on lifespan and error rate, but also on the absolute number of ALLD agents, which in this case scales with the size of the overall population. Figure 7(a) shows that direct observation is often the best strategy when there are fifty ALLD agents in a population of one hundred fifty. Even here, Gossipers can outperform Direct Observers if the accuracy of communication is sufficiently high and if the lifespan of the agents is fairly short. In contrast, Figure 7(b) illustrates that the Gossiping strategy is often preferred when there are five-hundred ALLD agents. Increasing the size of the population decreases the chance of the Direct Observer strategy prevailing. In contrast, the Direct Observer strategy does well when communication is highly erroneous or agents live a long time (cf. Equation 8).

Direct Observer success depends on the number of ALLD agents because for total knowledge they must individually experience each agent. As the size of the population grows, the Direct Observer must spend much longer to categorize the same fraction of agents. In contrast, the Gossip strategy can more easily scale to growing populations. Because the entire population works together to uncover defectors, their success is much less perturbed by population growth. In summary, we have shown that the success of gossiping compared to direct observation is dependent not only on lifespan, the rate of communication and its fidelity, but also on the size of the population.

3 General Discussion

We have replicated in a spatial model the Nakamura and Masuda (2011) results and confirmed their analysis, that gossip is able to sustain cooperation even when the information transmitted contains a high degree of uncertainty. We propose an explanation for this: when assessing the value of information in an evolutionary context we must incorporate not only accuracy but also relevance. We tested this in an experiment which examined how transmission rate interacts with transmission error. We found that socially-learned information even with a high incidence of error may outperform experience-based learning if the individuals whose reputation is learned by the latter method are not the same as those the agent subsequently encounters. In this we made the assumption that gossip is always as fast as or faster than individual learning.

Finally we have demonstrated that the relative importance of accuracy and relevance is highly sensitive to the ecology of the group, where by ‘ecology’ we mean demographic and life-history characteristics determining network diffusion properties. We do not argue that gossiping individuals will always outperform those that learn from direct experience. Our conclusion is that even if the incidence of error in gossip is high, which is a reasonable assumption given that gossip is second-hand information and talk is cheap, this does not mean that gossip is necessarily an inferior source of information.

A possible objection to this claim is that if changes in a society's ecology increase the frequency of re-meeting between individuals then the scope to offset the error in gossip by its superior relevance is reduced and learning by experience might prove more reliable (Roberts, 2008). Equation 8 shows that in fact either technique may be preferred, depending also on the context of transmission rate, transmission fidelity and lifespan. Gossip could still be a means of defence against rogue individuals that travel further. Parochialism might be another defence, but this would also limit opportunities for beneficial exchange between communities (Abbink et al., 2012; Rand et al., 2014).

The potential for increasing the probability of re-meetings will also vary according to the interaction in question. Some types of interaction naturally take place very often during an individual's lifespan, while others occur much less frequently. Hence as the results of Experiment 3 indicate, gossip may be more relevant for interactions that are infrequent in the context of an agent's lifespan.

Importantly, we have gone to lengths to give Direct Observation perhaps unreasonable levels of advantage in these simulations. In the case where perception of events was uncertain even for the immediate observer (as often occurs in human socialisation), then if transmission fidelity is higher than perception fidelity, it could be possible to *improve* the accuracy of information by combining information based on the perception of multiple sources (Fahle and Poggio, 1981; Moravec, 1988; Jacobs et al., 1991; Kao and Couzin, 2014). If for example learning by experience were to involve direct observation as opposed to having been directly involved in the exchange, then we are able to question the assumption that $k = 1$. This is true even for direct exchanges if the values of the costs and benefits are for any reason obscured to those involved, a condition common in both treaties and commerce.

Note that our results do not impact on the issues faced in explaining cooperation *per se*, for example the objection that gossip itself assumes cooperation and that therefore using it to explain cooperation is a regress. Of course, the invasion of individually maladaptive but socially beneficial traits has been well explored and established elsewhere (e.g. Čače and Bryson, 2007; Alizon and Taylor, 2008). However, our principle interest here is not so much to explain the ubiquitous presence of cooperation. Here we take gossip and cooperation as a setting in which we have addressed the claim that cheap, unreliable communication could not supplant more reliable direct observation.

4 Conclusion

Intuitively we tend to value accurate sources of information over inaccurate ones. Under the influence of this intuition the evolution of socially-acquired information seems hard to explain as its accuracy is likely to be lower than information acquired by direct experience. We have shown that the value of an information source cannot be determined by its accuracy alone, relevance must also be taken into account. Once this is done it is possible to explain that despite even a high level of inaccuracy, an individual's best information-gathering strategy may still be to depend on gossip.

5 Acknowledgments

We would like to thank Gilbert Roberts for helpful discussion regarding the issue of frequency of interactions, and Willem Frankenhuis and our anonymous reviewers for helpful comments and suggestions.

References

Abbink, K., Brandts, J., Herrmann, B., and Orzen, H. (2012). Parochial altruism in inter-group conflicts. *Economics Letters*, 117(1):45–48.

Alexander, R. D. (1987). *The Biology of Moral Systems.* Foundations of Human Behavior. Aldine de Gruyter, New York.

Alizon, S. and Taylor, P. (2008). Empty sites can promote altruistic behaviour. *Evolution*, 62(6):1335–1344.

Asch, S. E. (1956). Studies of independence and conformity: I. a minority of one against a unanimous majority. *Psychological Monographs: General and Applied*, 70(9):1–70.

Axelrod, R. and Hamilton, W. D. (1981). The evolution of cooperation. *Science*, 211:1390–1396.

Dawkins, R. (1976). Hierarchical organisation: A candidate principle for ethology. In Bateson, P. P. G. and Hinde, R. A., editors, *Growing Points in Ethology*, pages 7–54. Cambridge University Press, Cambridge.

Dessalles, J.-L. (2007). *Why we talk: the evolutionary origins of language.* Oxford University Press.

Enquist, M. and Leimar, O. (1993). The evolution of cooperation in mobile organisms. *Animal Behaviour*, 45:747–757.

Fahle, M. and Poggio, T. (1981). Visual hyperacuity: Spatiotemporal interpolation in human vision. *Proceedings of the Royal Society of London B: Biological Sciences*, 213(1193):451–477.

Gergely, G., Bekkering, H., and Király, I. (2002). Rational imitation in preverbal infants. *Nature*, 415:755.

Gergely, G. and Csibra, G. (2006). *Sylvia's recipe: the role of imitation and pedagogy in the transmission of cultural knowledge*, pages 229–255. Roots of Human Sociality: Culture, Cognition and Interaction. Berg, New York.

Goldberg, D. and Deb, K. (1991). A comparative analysis of selection schemes used in genetic algorithms. *Urbana*, 51:61801–2996.

Gruber, T., Poisot, T., Zuberbühler, K., Hoppitt, W., and Hobaiter, C. (2015). The spread of a novel behavior in wild chimpanzees: New insights into the ape cultural mind. *Communicative & Integrative Biology*, 8(2):e1017164.

Hardin, G. (1968). The tragedy of the commons. *Science*, 162:1243–1248.

Henrich, J. and McElreath, R. (2003). The evolution of cultural evolution. *Evolutionary Anthropology: Issues, News, and Reviews*, 12(3):123–135.

Hess, N. and Hagen, E. (2006). Psychological adaptations for assessing gossip veracity. *Human Nature*, 17(3):337–354.

Jacobs, R. A., Jordan, M. I., Nowlan, S. J., and Hinton, G. E. (1991). Adaptive mixtures of local experts. *Neural computation*, 3(1):79–87.

Kao, A. B. and Couzin, I. D. (2014). Decision accuracy in complex environments is often maximized by small group sizes. *Proceedings of the Royal Society B: Biological Sciences*, 281(1784).

Kokko, H. (1997). Evolutionarily stable strategies of age-dependent sexual advertisement. *Behavioral Ecology and Sociobiology*, 41(2):99–107.

McElreath, R. and Boyd, R. (2007). *Mathematical models of social evolution: A guide for the perplexed.* University of Chicago Press.

Meltzoff, A. N. (1988). Infant imitation after a 1-week delay: Long-term memory for novel acts and multiple stimuli. *Developmental psychology*, 24(4):470.

Milinski, M., Semmann, D., and Krambeck, H.-J. (2006). Reputation helps solve the tragedy of the commons. *Nature*, 415(6870).

Moravec, H. P. (1988). Sensor fusion in certainty grids for mobile robots. *AI magazine*, 9(2):61.

Nagell, K., Olguin, R. S., and Tomasello, M. (1993). Processes of social learning in the tool use of chimpanzees (pan troglodytes) and human children (homo sapiens). *Journal of Comparative Psychology*, 107(2):174–186.

Nakamura, M. and Masuda, N. (2011). Indirect reciprocity under incomplete observation. *PLoS Computational Biology*, 7(7):e1002113.

Nowak, M. A. and Sigmund, K. (1998). Evolution of indirect reciprocity by image scoring. *Nature*, 393:573–577.

Nowak, M. A. and Sigmund, K. (2005). Evolution of indirect reciprocity. *Nature*, 437:1291–1298.

Ohtsuki, H. and Iwasa, Y. (2006). The leading eight: Social norms that can maintain cooperation by indirect reciprocity. *Journal of Theoretical Biology*, 239(4):435 – 444.

Ohtsuki, H., Iwasa, Y., and Nowak, M. A. (2009). Indirect reciprocity provides only a narrow margin of efficiency for costly punishment. *Nature*, 457(7225):79–82.

Rand, D. G., Nowak, M. A., Fowler, J. H., and Christakis, N. A. (2014). Static network structure can stabilize human cooperation. *Proceedings of the National Academy of Sciences*, 111(48):17093–17098.

Rauwolf, P., Mitchell, D., and Bryson, J. J. (2015). Value homophily benefits cooperation but motivates employing incorrect social information. *Journal of Theoretical Biology*, 367(0):246–261.

Richerson, P. J. and Boyd, R. (2005). *Not by Genes Alone: How Culture Transformed Human Evolution*. Chicago University Press, Chicago.

Roberts, G. (2008). Evolution of direct and indirect reciprocity. *Proceedings of the Royal Society B: Biological Sciences*, 275(1631):173–179.

Schelling, T. (1960). *The strategy of conflict*. Harvard University Press.

Scott-Phillips, T. C. (2008). Defining biological communication. *Journal of Evolutionary Biology*, 21(2):387–395.

Sommerfeld, R. D., Krambeck, H. J., Semmann, D., and Milinski, M. (2007). Gossip as an alternative for direct observation in games of indirect reciprocity. *Proceedings of the National Academy of Sciences of the United States of America*, 104(44):17435–17440. PUBM: Print-Electronic; DEP: 20071018; JID: 7505876; 2007/10/18 [aheadofprint]; ppublish.

Sperber, D. and Wilson, D. (1995). *Relevance: Communication and Cognition*. Wiley.

Taylor, D. J. (2014). *Evolution of the Social Contract*. PhD thesis, Department of Computer Science, University of Bath.

Tomasello, M. (1999). *The Cultural Origins of Human Cognition*. Harvard University Press, Cambridge, MA.

Trivers, R. L. (1971). The evolution of reciprocal altruism. *Quarterly Review of Biology*, 46(1):35–57.

Čače, I. and Bryson, J. J. (2007). Agent based modelling of communication costs: Why information can be free. In Lyon, C., Nehaniv, C. L., and Cangelosi, A., editors, *Emergence and Evolution of Linguistic Communication*, pages 305–322. Springer, London.

Whiten, A., McGuigan, N., Marshall-Pescini, S., and Hopper, L. M. (2009). Emulation, imitation, over-imitation and the scope of culture for child and chimpanzee. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 364(1528):2417–2428.

Williams, G. C. (1966). *Adaptation and Natural Selection*. Princeton University Press.

Yoo, J. H. (2009). The power of sharing negative information in a dyadic context. *Communication Reports*, 22(1):29–40.