

The Cooperative Origins of Epistemic Rationality?

Corey Dethier

[Draft; forthcoming in *Erkenntnis*]

Abstract

Recently, both evolutionary anthropologists and some philosophers have argued that cooperative social settings unique to humans play an important role in the development of both our cognitive capacities and the “construction” of “normative rationality” or “a normative point of view as a self-regulating mechanism” (Tomasello 2017, 38). In this article, I use evolutionary game theory to evaluate the plausibility of the claim that cooperation fosters epistemic rationality. Employing an extension of signal-receiver games that I term “telephone games,” I show that cooperative contexts work as advertised: under plausible conditions, these scenarios favor epistemically rational agents over irrational ones designed to do just as well as them in non-interactive contexts. I then show that the basic results are strengthened by introducing complications that make the game more realistic.

0 Introduction

Why shouldn’t I believe whatever it will make me happiest to believe? A common answer is that this simply isn’t how beliefs work: you don’t get to choose whether or not you believe something (Hieronimi 2005). This answer raises a new question: why not? Why do we feel compelled to believe what it is “rational” to believe, and from where does the deep normative hold of this epistemic notion arise? Recently, both evolutionary anthropologists (Mercier and Sperber 2017; Tomasello 2014, 2017, 2020) and some philosophers (Dogramaci 2012, 2015; Graham 2020; Sharadin 2018) have argued for a particular evolutionary answer to this set of questions: the kinds of cooperation unique to human testimonial practices promote the development of both our unique cognitive abilities and the “construction” of “normative rationality” or “a normative point of view as a self-regulating mechanism, arguably the capstone of the ontogeny of uniquely human cognition” (Tomasello 2017, 38). Call this the “cooperative hypothesis.”

In this article, I focus my attention on the second aspect of the evolutionary story outlined above, namely the claim that cooperation fosters epistemic rationality. Employing an extension of signal-receiver games that I term “telephone games,” I show that cooperative contexts work as advertised: under plausible conditions, these scenarios favor epistemically rational agents over

irrational ones designed to do just as well as them in non-cooperative contexts. I then show that the basic results are strengthened by introducing complications that make the game more realistic.

I begin by outlining the cooperative hypothesis, particularly as defended by Tomasello (§1). I then define a class of cooperative interactions that I call “telephone games,” the crucial feature of which is that they allow for variation in the number of agents that serve as intermediaries between an observation and an action, and lay out two different strategies—one epistemically rational and one “wishful”—for playing the game (§2). I show that while wishful agents do outperform rational ones under some cooperative conditions, rational agents are generally favored when the initial signal is at least moderately reliable and group size is three or more (§3). I then introduce a complication, embedding the telephone game within a stag-hunt-style scenario that indicates that rational agents are usually favored when cooperation is (§4). The results outlined in §3 and §4 offer support to the cooperative hypothesis, though they have limitations, as I discuss in the final section (§5). Throughout, my presentation will be largely informal; technical details and results can be found in the appendix.

1 Background: the cooperative hypothesis

Whatever one’s preferred account of rationality—and here I include accounts in which rationality is “bounded” (Gigerenzer 2008) or context-sensitive (Morton 2017)—I take it that humans are not actually, or at least not always, rational. That’s true regardless of whether the notion of rationality at issue is practical or epistemic: there’s disagreement about how well we approximate and how frequently we deviate from these ideal notions, but there is broad agreement that humans are not perfectly rational. Nevertheless, we generally take the rules of rationality to be normatively compelling even though we consistently fail to live up to them: while it’s forgivable and understandable for an agent to fail to be perfectly epistemically rational, one is performing better epistemically the closer one is to approximating perfectly rational behavior. Finally, despite the fact that we deviate from them *all the time*, the norms of epistemic rationality are so deeply ingrained in our thinking that we typically find ourselves effectively unable to consciously violate them, or at least some of them. It’s commonly accepted in epistemology, for example, that we can’t bring ourselves to believe something simply because we want it to be true (see, e.g., Hieronymi 2005). All of this is purely

descriptive, and so, presumably, all of the above features of our epistemic life are open to relatively straightforward evolutionary explanations.

At least in broad strokes, the imperfect nature of our reasoning faculties is the easiest to explain evolutionarily (though of course there remains significant disagreement about the details): supposing that a perfect reasoning system is desirable but extremely “costly,” we’re unlikely to develop perfect reasoning if cheaper approximations and heuristics work almost as well in the most evolutionarily important contexts. More puzzling is why we should recognize and take ourselves to be bound by norms of epistemic rationality—norms that deviate from our actual behavior—and why these norms should be so deeply ingrained that we don’t just have intuitions about their correctness but often find ourselves incapable of consciously violating them. Recently, a number of evolutionary anthropologists and philosophers have suggested that our penchant for cooperation provides the answer to this question: we recognize and obey these norms because (approximating) epistemic rationality plays an important role in facilitating cooperation.¹ Epistemic norms have a hold on us, in other words, because we are deeply cooperative in a way that our closest relatives aren’t.

Tomasello (2014, 2017, 2020) and his collaborators (e.g., Tomasello et al. 2012) have gathered substantial experimental evidence in support of the more general claim that human cognitive faculties are evolutionarily linked to cooperation. Humans have a proclivity to cooperate and skills (or perhaps instincts) that facilitate cooperation and that are lacking in our closest relatives. Both proclivity and skills develop relatively early and their development is connected, both in time and in ability, with the development of the unique features of human linguistic competency. Notably, while the skills necessary for competitive interaction are also clearly linked to other aspects of our cognitive development, they don’t seem to be linked to linguistic competency and, further, they’re largely shared with our closest relatives. The inferred conclusion: the cognitive differences that separate us from our closest relatives exist because of the evolutionary pressures of cooperation, particularly cooperation involving the (linguistic) exchange of information.

One of these cognitive differences—as Tomasello discusses—is that humans recognize, and take themselves to be normatively constrained by, an epistemic notion of rationality. From a young age,

¹There are some noteworthy differences between the treatment given by anthropologists and that given by philosophers. Tomasello (2014) and Mercier and Sperber (2017) are primarily focused on why we recognize epistemic norms. Dogramaci (2012, 2015), Graham (2020), and Sharadin (2018) are more concerned with normative questions. My focus in this essay will be on the former.

humans are deeply concerned not just with what is believed but also with what ought to be. This normative component has no obvious analogue in our nearest relatives, and Tomasello argues that cooperation is responsible for this specific aspect of our cognitive life in two ways. First, recognizing an epistemic norm as such requires both concepts and cognitive skills that seem to be lacking in our closest relatives. For instance: epistemic norms are only *epistemic* insofar as they relate to the truth or accuracy of mental representations, but the notion of “truth,” or more precisely of a sort of objective perspective that may depart from my own, is itself a substantial cognitive achievement—and one that Tomasello shows that children make use of in their cooperative interactions but great apes seems incapable of recognizing (Tomasello 2014, 38-45). Similar comments apply to the ability to “self-monitor” that’s presupposed in taking an epistemic norm to be genuinely *normative* rather than merely habitual (Tomasello 2014, 58). Second, in addition to providing humans with the skills and concepts to “construct” this “normative point of view” (Tomasello 2017, 38), cooperation also provides the motivation for doing so: in order to attract new partners for future cooperative tasks, we must make it clear that our testimony reflects the truth rather than our own selfish interests (Tomasello 2014, 75). Clear obedience to shared epistemic norms marks a potential cooperative partner as one whose testimony will be reliable. As Tomasello tells the story, this need to be recognized as a reliable partner eventually becomes a deeply ingrained social institution, a set of (epistemic) norms that we find ourselves incapable of consciously deviating from. This account of the origin of epistemic norms fits with both his broader picture of the development of human normative judgments (Tomasello 2016) and the more general literature on the latter subject (Bicchieri 2006; Sterelny 2012).

For Tomasello, therefore, epistemic rationality involves being normatively guided or constrained by an objective or unbiased perspective towards evidence. That is: epistemic rationality involves abiding by objectivity norms that prescribe against biased evaluation of the evidence. What I’m calling “the cooperative hypothesis” is his claim that cooperative scenarios fostered the development of epistemic rationality in this sense: humans abide by this kind of prescription for objectivity *because* it is advantageous to do so in cooperative contexts. Though there are substantial disagreements about the nature of epistemic rationality, I take it that the conception of rationality laid out by Tomasello is relatively widely shared in that these disagreements can be understood as concerning the nature of the “objective” or “unbiased” perspective or the kind of normative constraint

that such norms offer.² We might therefore think of Tomasello’s conception of epistemic rationality as a “minimal” one: it’s a characterization that fits a variety of accounts that differ substantively in their details. And the cooperative hypothesis can thus be understood as the claim that *whatever* the specific details of “objectivity,” we have norms that prescribe it because such norms are advantageous in cooperative contexts.

Philosophers and evolutionary game theorists have long been interested in human cooperation (see Bicchieri 1997, 2006; Skyrms 1996, 2004), and there’s a growing literature on both the development of particularly epistemic cooperation (Zollman 2017) and the fixation of the human instinct for cooperation (Bicchieri 2006; Forber and Smead 2015; McLoone and Smead 2014; McLoone et al. 2018). The aim of investigating human cooperation using tools from evolutionary game theory is to determine whether or evolutionary explanations like the one offered by Tomasello are plausible: is it the case that cooperative settings actually promote epistemic rationality? If such scenarios do not promote rationality in the controlled models of evolutionary game theory, then the cooperative hypothesis is unlikely to be true without substantive amendment. By contrast, if such scenarios do promote rationality, then we’ve shown that the mechanism proposed by Tomasello can work, though (obviously) more work is needed to show that cooperation is in fact responsible for the historical development of epistemic norms.

2 Telephone games

Imagine an agent who observes their environment and then interacts with it on the basis of the information they gain from the observation. In such a scenario, there are zero degrees of separation between observation and action. Increase the degrees of separation and the scenario comes to be similar to the children’s game “telephone”: one agent provides the original message, and it is passed from intermediary to intermediary until it reaches the actor—hopefully without distortion. Suppose that we attach payoffs to this scenario: the final agent in the chain picks one of a group of possible actions and if their choice is the correct one, everyone shares in the payoff. Call games with this structure “telephone games”; for all intents and purposes, telephone games are simply an extension of traditional “sender-receiver” games introduced by Lewis (1969) and popularized by Skyrms (see

²This is true even of some (but not all) of those philosophers suspicious of traditional understandings of objectivity; Douglas (see, e.g., 2009).

	s_1	s_2	s_3
a_1	2	0	0
a_2	0	1	0
a_3	0	0	1

	s_1	s_2	s_3
a_1	1	0	0
a_2	0	2	0
a_3	0	0	1

	s_1	s_2	s_3
a_1	1	0	0
a_2	0	1	0
a_3	0	0	2

Fig. 1: Three possible payoff matrices for a game with three states and three corresponding actions.

Skyrms 2010).³

The crucial features of telephone games are that they are cooperative and that they involve testimony. The cooperative element is built into the payoffs of the game: each agent receives a payoff if and only if the agent at the end of the chain of telephone picks the correct action; one agent profits if and only if all agents do. I’ll allow agents to differ in their preferences, however (see figure 1). The effect of this stipulation is that agents with less-than decisive evidence can differ in their choices: an agent who profits much more in one scenario than in the other may try for the first option even if their evidence indicates that the second option is more likely to be successful. The testimonial element, in turn, introduces a natural means of distinguishing between what the agent believes to be true from their dispositions to act. This distinction allows us to separate genuinely wishful thinkers, who believe what they want to be true, from agents who merely take their preferences into account when acting. There’s nothing wrong with the latter. In fact, it’s what a utility-maximizing agent is supposed to do. By contrast, there’s supposed to be something wrong—epistemic irrationality—with taking preferences into account when believing or testifying on the basis of belief. If this normative intuition has an evolutionary basis, then agents that keep separate their beliefs and desires should outperform agents that don’t in at least some scenarios.

Slightly more technically, telephone games consist of the following. A state is picked from among S equiprobable and mutually exclusive alternatives, and the first agent in the game makes an observation, o . They are assumed to know the objective probability of each state given this observation; that is, $Pr(s_i|o)$ for all $s_i \in S$.⁴ In what follows, I’ll talk about the mean “power” of the original

³As one reviewer helpfully reminded me, the game is often called “Chinese whispers,” and there’s a fairly long history of analogies between evolutionary mechanisms of various sorts and “Chinese whispers” (see, e.g., Ridley 2001). I have yet to encounter a formalization using game theoretic tools along the lines developed here, but it would not surprise me to discover one.

⁴As noted above, there are interesting debates about how to understand the sense of objectivity involved in epistemic rationality, and there may well be cases in which there’s no such thing as an objective probability that might complicate the analysis. In such cases, it looks like the proposed epistemic norms suffer a kind of presupposition failure. For present purposes, however, I’ll assume that the cases of interest are ones in which there are objective probabilities in the form of well-defined long-run frequencies. The question is whether hewing to such probabilities is evolutionarily advantageous.




Fig. 2: Two instances of a telephone game. The left-hand game is a 1-degree game with a single wishful observer in addition to the actor (whose type is irrelevant). The right-hand game is an instance of a 3-degree game with two wishful agents and one rational one.

observation, by which I intend the mean probability of the true state of affairs given the observation, or $Pr(s_t|o)$.⁵ The aim of the game is to pick the action that “matches” the true state of the world, though agents will always prefer some matches to others; which of the states an agent prefers is picked randomly from the available states. A signal is then passed from the initial observer to another agent n times, where n is what I earlier termed the “degree” of the telephone game—i.e., in a zero-degree game, the observer and actor are the same agent; in a one-degree game, the observer passes the signal directly to the actor, etc. Supposing that all the agents are “honest,” this signal will simply consist in the probability that they assign to each of the possible states s_i .

I allow agents to differ solely in how they form their subjective beliefs. What I call “rational” agents take both observations and testimony at face value; if their information indicates that the objective probability of s_i is x , then they assign probability of x to s_i . Their own preferences—i.e., that they receive a larger payoff for s_j than s_i —play no role in what probabilities they assign. They are, in other words, “objective” or “unbiased” in their evaluation of the evidence; they’re epistemically rational in the minimal sense discussed in the last section. By contrast, the agents that I term “wishful” take their preferences into account when assigning probabilities to various states. So, for example, if they prefer s_j to s_i , they’ll assign a higher probability to s_j than it objectively has on the evidence available. They are, in other words, “wishful thinkers”: they believe what they want to be true. Since our agents are all honest, they all testify in accordance with their beliefs, and the effect is that rational and wishful agents have different patterns of testimony: the rational agent always passes on the information that they received without loss, while the wishful agent introduces some (self-serving) noise.⁶ (For technical details on how all this works, see the appendix.)

⁵In all of the simulations conducted in this paper, the standard deviation for observation power was .05. Test runs with other amounts of variance indicated that, as expected, increasing / decreasing the variance simply served to smooth / roughen the observed effect.

⁶Note that we could interpret this game differently: we could understand the rational agents as honest and the wishful ones as dishonest (in a particular way). That is: in cooperative scenarios, wishful thinking can be seen as akin to a kind of dishonesty.

Since the goal is to isolate the effect of cooperation of belief, we want the two different types of actors to behave the same way outside of such contexts. This is easily accomplished by stipulating that the wishful agents choose which action to take based purely on what they believe (i.e., they don't double-count their preferences). This stipulation guarantees that the two strategies will perform identically in non-interactive degree-0 telephone games: both will act to maximize their expected utility based on a signal that they get directly from the world (compare Okasha 2013, 753). The steps they take in calculating the expected utility will be different, but the outcome will be the same. The upshot is that we can use these two strategies to isolate the effects of introducing cooperation: since they perform identically in the degree-0 scenario, any difference in their performance in higher-degree scenarios must be attributed to the increase in the number of agents involved in the game.

Notice one further point. Wishful agents testify as they would act, rather than as they believe. Many types of epistemic irrationality look like they're guaranteed to fail because they'll lead agents to miscalculate expected utilities and be instrumentally worse off (Okasha 2013, 2018). By contrast, the kind of constrained wishful thinking defined here looks like it should *outperform* epistemic rationality if anything: the wishful agent chooses the belief state with the highest expected utility, whereas the rational agent merely chooses the one that is most likely to be true.⁷ As such, it should be (mildly) surprising if rational agents were employing the optimal or even preferable strategy. And thus if we observe that result, it serves as relatively powerful evidence for the cooperative hypothesis.

The results that follow were generated using a standard discrete version of the replicator dynamics.⁸ In slightly more detail, I assumed that both strategies were present in an infinite population and freely-mixing—i.e., there was no preference among rational or wishful agents to play with either rational or wishful agents. Given a proportion of prevalence for the two strategies, each strategy

⁷Well, ok, that's not *quite* right: essentially, the wishful agent maximizes the expected utility of their beliefs relative to a restricted subset of possible outcomes—namely, non-cooperative scenarios. One way of viewing the results that follow in this paper is as showing that this restricted subset is *not* representative: an agent well-designed for non-cooperative scenarios is not necessarily well-designed for cooperative ones. Even the wishful thinkers that I've designed here, who are unrealistically good at managing traditional problems with overconfidence in the setting of individual belief-formation, can have trouble in social contexts, because the latter present pitfalls for wishful thinkers that are unpredictable from a purely individualist perspective.

⁸For discussion of the modeling assumptions involved in the use of the replicator dynamics, see O'Connor (2020). For discussion of the mathematical differences between different versions of the discrete replicator dynamics, see Pandit, Mukhopadhyay, and Chakraborty (2018).

has an expected payoff determined by how likely they are to interact with the other strategy. If the rational strategy has a higher payoff than the wishful one, then the proportion of rational to wishful agents increases according to how substantial the difference in payoff is. Eventually, one strategy takes over or the two reach the point where neither is advantaged. (Again, details can be found in the appendix.)

Before touching on the results themselves, it’s worth noting a limitation of telephone games as just outlined. As I’ve set things up, payoffs are uniform across a strategy and agents have no choice but to involve themselves in the practice of providing and (more importantly) trusting testimony. Essentially, these restrictions amount to idealizing away any “undermining” behaviors of the agents: they’re not allowed to cheat or opt-out of cooperating.⁹ I’ll relax this restriction somewhat below, where I’ll consider games in which agents can opt-out of cooperation, and I’ll address the limitations introduced by these idealizations at more length in §5.

3 Results

The results are as follows. For any game of degree 2 or more (recall: involving at least three agents in the chain), there is a minimal level of mean observation power above which every run of the model results in $> 99\%$ of all agents playing the rational strategy. Changing the degree of the system has only a very minimal effect on this minimal power level, while increasing the cardinality substantially lowers it. In any degree 2 or more game, though, rational agents are favored so long as the average power of the observation is $> 65\%$ —and are often favored in situations where power is substantially less than that. The results of degree 1—that is, 2-player—games are a bit more bizarre, with alternating areas favoring both strategies and—as cardinality increases—large regions in which neither strategy outperforms the other. Figure 3 contrasts degree-1 games with games of higher degrees. While I’ve reproduced results for a wide range of power values, it should be noted that values at the low end of the scale are in an important sense less realistic: given that the relevant notion of “observation power” is the objective probability of the state that actually obtains on the

⁹As one reviewer rightly noted, we should expect that agents in this kind of cooperative situation would develop mechanisms to determine what kind of agent was providing them with testimony. Since these would essentially be cheater-detection mechanisms, I take it that their development and operation will introduce a familiar arms-race between detectors and cheaters. While it’s an interesting question whether this arms-race will play out in the same way as in more familiar scenarios, issues of space require that question to be left to future work. (See also the end of §4).

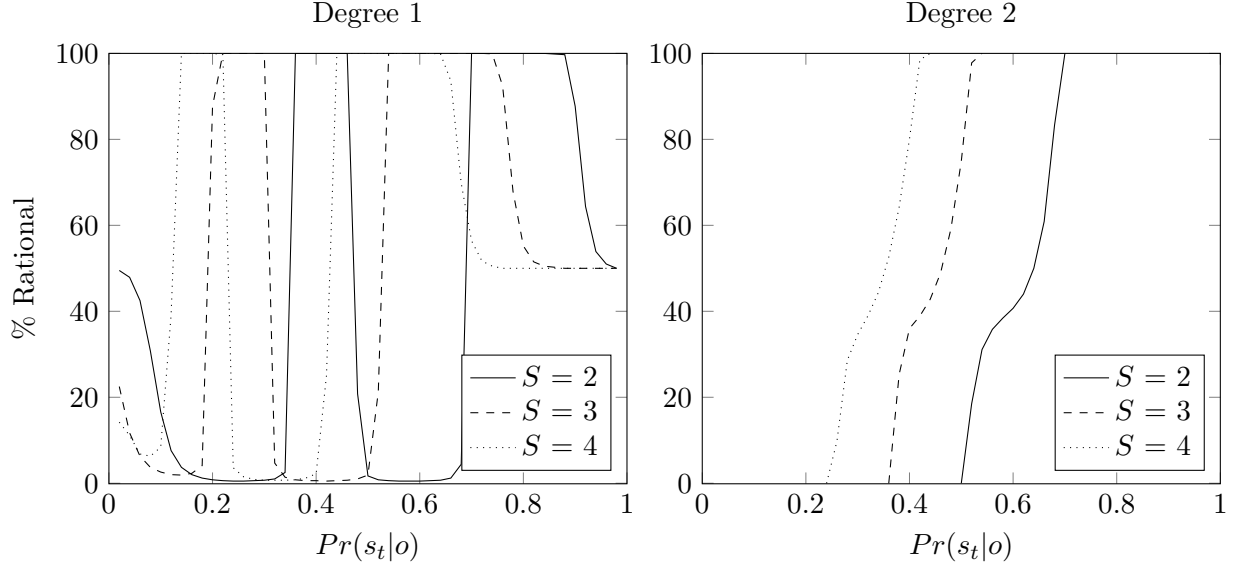


Fig. 3: The percentage of agents playing the rational strategy when the game halts (% Rational) given different mean observation powers ($Pr(s_t|o)$) and cardinalities (S). Degrees 3 and 4 are largely identical to degree 2. Standard deviation for observation power is .05 throughout.

evidence ($Pr(s_t|o)$), the probability of observations with low power must be assumed to be low.

Mathematically speaking, there are a few different factors that explain how the model generates these results. The most important factor is the following. For a game of given degree, cardinality, and power, there’s a set degree of “noise”—in the sense of agential preferences—that must be incorporated into the signal before the actor will chose a state other than the one that the signal gives the highest probability to. In cases in which the signal is highly misleading, this noise is beneficial because it can lead to agents picking the action that matches the true state even though the evidence indicates that that state doesn’t obtain; this is why wishful agents are generally favored for low values of $Pr(s_t|o)$. When the signal is accurate, however, the noise introduced by wishful agents becomes disadvantageous, and for the same reason: it can lead the actor to ignore the evidence, this time to everyone’s detriment. This basic fact about when it’s good to add noise explains the general correlation between higher values of $Pr(s_t|o)$ and rational agents being favored found throughout games of all degree and cardinalities.

The rational strategy is most favored when the signal is accurate but not too accurate; the number of turns required for the rational strategy to take over increases as $Pr(s_t|o)$ approaches 1. The closer the observation power is to 1 (or 0) the more noise will be required for the chain of agents

to end up making a choice different from the one with the highest probability on the evidence. What happens in the model is that as $Pr(s_t|o)$ approaches 1, chains incorporating wishful agents choose an action that differs from that chosen by a chain of only rational agents with decreasing frequency; an increased amount of noise is required to obscure the evidence. In degree-1 games (particularly at higher cardinalities), high values for observation power are equally favorable for both agents because wishful agents cannot introduce enough noise to lead the actor away from picking the state that the evidence virtually guarantees is the correct one.

The odd behavior of degree-1 games when compared to higher-degree games is partly explained by the fact that wishful agents “catch up” to rational ones faster in these games and partly by a second phenomenon: there are values of $Pr(s_t|o)$ for which adding a wishful observer can cause the decision to be *more* responsive to the evidence than it is when the observer is rational. So consider a situation in which a rational observer will report that one of two states is 1.5 times as likely as the other. If the actor prefers one of the states by a ratio to 2:1, then the evidence is irrelevant: the actor will always choose their preferred state. As we would expect, this insensitivity to the evidence is on average good when the evidence is relatively misleading and bad when it’s more reliable. Switching the rational observer for a wishful one changes the reported probabilities, and can thus have the effect of changing the situation from one in which the evidence is irrelevant to one in which it matters. Whether or not the wishful agent’s preferences have this effect depends on whether their preferences align with those of the actor. If they both prefer the non-evidenced state, the evidence will be ignored; in all other cases, the actor will pick whatever state the evidence suggests is the correct one. In the case where the cardinality is 2 and the degree is 1, the result is that ignoring the evidence is twice as likely when paired with a rational agent than when paired with a wishful one. The same phenomenon occurs at higher degrees, but the difference in the probability of observing the effect decreases rapidly with degree, rendering the phenomenon essentially irrelevant above degree 1.

4 Introducing complications

In the telephone games presented above, actors have no choice but to cooperate: the game is simply one in which the payoffs are cooperative. Such games are hardly common, however; even

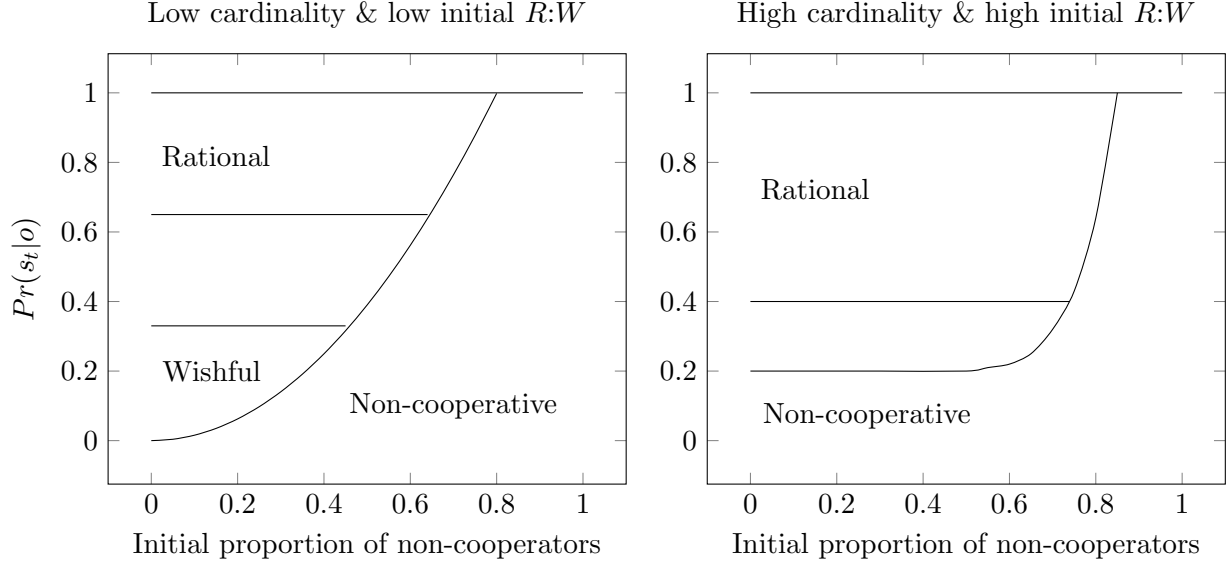


Fig. 4: Schematic representation of the conditions under which each of the three strategies goes to fixation in different games. The left-hand scenario is representative of games of low cardinality and low initial proportion of cooperators playing the rational strategy; the right hand, games of high cardinality and high initial proportion of rational players. As can be seen, tuning these values increases the performance of rational agents relative to both wishful agents and non-cooperators, though it also improves the performance of non-cooperators relative to wishful ones. The unmarked area represents conditions that result in a mix of rational and wishful agents.

in cooperative scenarios, there are aspects of competition and mechanisms to prevent cheating are required to ensure that the payoffs remain cooperative (Forber and Smead 2015; c.f. Bicchieri 2006). Further, though Tomasello downplays these mechanisms relative to other authors, he also quite clearly sees them as important, stressing the role of partner choice—and to a lesser extent, partner control—in ensuring that agents behave rationally (see, e.g. Tomasello 2014, 75, 89).

Better capturing these aspects of the purported evolutionary scenario requires complicating the game. More specifically, I embedded the basic telephone game within a stag-hunt-like scenario (see Skyrms 2004). Essentially, actors have two choices: risk playing the telephone game, or “hunt hare” by taking a guaranteed payoff. Adding this element introduces an additional variable into our game, namely the percentage of cooperators in the system. All other things being equal, the more cooperators the less risky it is to be a cooperator. The expectation: if the telephone game pays badly enough—because the signal is too poor, for example—then cooperators will suffer and eventually die out.

Figure 4 presents a schematic representation of the results (a more technical discussion can

be found in the appendix). As in the basic telephone game, the primary variable determining whether rational or wishful agents thrived is the strength of the signal, with rational agents outperforming wishful agents at higher signal strengths. As expected, the primary factor affecting whether non-cooperators outperformed cooperators is the initial proportion of cooperators to non-cooperators, though signal strength plays an important role here as well, with non-cooperators competing with wishful agents for dominance in low-power settings. Increasing the cardinality, the initial percentage of cooperators playing rational, or both has the effect of increasing the advantage for non-cooperators at low power values while decreasing their effectiveness in high-power settings—effectively, tuning these variables high enough ensures that any situation in which the wishful strategy is favored is one in which non-cooperation is favored more.

In all runs of the game, a high enough initial proportion of non-cooperators ($\geq 85\%$) was sufficient to prevent any cooperative strategy from going to fixation, no matter how the other variables were tuned—though some test runs with more extreme advantages for cooperators did show the ability to push the required initial proportion up to close to 95%. As with standard stag hunt games, therefore, there’s a problem with getting cooperation off the ground. To better capture how the type of cooperative testimony represented in telephone games might get started, therefore, we’d probably need to introduce more structure into our population in the form of recognition mechanisms or other means of ensuring that cooperators are more likely to play with their own kind. Since this is a familiar issue in stag hunt games, however, I’ll assume that standard complications have similar effects here as well.

5 Interpreting the results

The hypothesis that I set out to test holds that cooperative interaction between agents plays a central role in explaining our recognition of norms of epistemic rationality—specifically, in explaining why we recognize and obey objectivity norms. The rough idea is that cooperation introduces a pressure to be reliable—and not just reliable, but reliable in a way that will be apparent to other agents—in evaluating and representing evidence that isn’t present in non-interactive scenarios.

To test this hypothesis, I compared the performance of two agents who behave identically in non-interactive scenarios but who form mental representations and thus testify in different ways. The

epistemically rational agents testify to what their evidence indicates is the truth (i.e., they take what Tomasello refers to as an “objective” perspective on their evidence), while wishful agents testify based partly on what they want to be true (i.e., their testimony is based on features specific to their perspective). Generally speaking, the results favor the cooperative hypothesis: what the hypothesis predicts is that cooperative scenarios should favor obedience to epistemic norms; what we’ve seen is that—assuming that evidence isn’t too consistently misleading—cooperative scenarios do generally favor obedience to epistemic norms. As I noted above, this limiting assumption is reasonable, at least in the context of the game: given how I defined the power of the observation, we shouldn’t expect observations with extremely low power values to be particularly frequent. Increasing the complexity of the problems faced by agents—represented by the number of potential answers, i.e. the cardinality of the game—increases this advantage.

In §4, I showed that introducing non-cooperative agents into the mix has the effect of further disadvantaging wishful agents because non-cooperators generally outperform wishful thinkers where the latter are favored over rational agents. These results lend some support to the transcendental flavor of Tomasello’s claim that testimony isn’t possible outside of the context of epistemic norms that ensure reliability (Tomasello 2014, 51); more importantly, they help to demonstrate that situations that favor cooperative information-sharing broadly speaking will also usually favor epistemically rational information-sharing (see table 1). This last point provides what I take to be the strongest piece of evidence that I offer: given that we are extremely cooperative information-sharers, my results provide an indication that the scenarios that led us to be cooperative would also have led us to be epistemically rational. That’s not to say that the model shows that epistemic rationality *is* the result of selective pressures in cooperative scenarios; models don’t usually provide that sort of evidence (Forber 2010). Instead, the models indicate that the route to the development of epistemic rationality proposed by the cooperative hypothesis is possible. Given the backdrop of Tomasello’s extensive empirical evidence for the evolutionary connection between cooperation and cognitive faculties including epistemic norms, this fact should be seen as lending support to not just the cooperative hypothesis, but his overall account.

Still, there are two respects in which the results above don’t support Tomasello’s account. First, the contrast that I’ve tested is between non-interactive scenarios and cooperative ones; I haven’t tested the contrast between cooperative and competitive scenarios. So my results don’t support

	Simple model	Complex model			
		Low cardinality and R		High cardinality and R	
		Low N	High N	Low N	High N
Low power	wishful	wishful	non-coop	non-coop	non-coop
Med. power	mixed	mixed	non-coop	rational	rational
High power	rational	rational	non-coop	rational	rational

Table 1: Summary of which strategy is successful under the different conditions in each game.

Tomasello’s claim that it is our cooperative interactions *as opposed to* competitive interactions that foster epistemic rationality. I don’t think this seriously undermines the results, given my aim of testing whether the mechanism proposed by Tomasello can succeed—we’ve seen that it can.

The second of the key ways in which my model does not support Tomasello’s account concerns his reasoning for why rationality is supposed to be advantageous, namely that displaying rationality makes one an attractive partner for future cooperative interactions and enforcing it in partners ensures that their testimony is reliable. Indeed, this aspect of the account is crucial, as it’s what accounts for why epistemic norms are felt to be norms rather than mere habits. My model doesn’t respect this feature of the cooperative hypothesis, however: it’s not the case that rational agents are favored in telephone games because they are more likely to be chosen as partners. On the contrary, they’re favored in telephone games because they’re simply more likely to (be in chains that) pick the correct action to take in a given scenario. In my model, the norms of epistemic rationality are not norms at all, properly-speaking (compare Bicchieri 2006, 2017). One is simply better off being rational than not.

It would be natural, therefore, to read my results as showing that the social mechanisms postulated by the cooperative hypothesis are not necessary for a cooperative explanation of the existence of epistemic rationality. That would be a mistake, however. The better way to view my results is as agnostic as to the mechanism by which cooperative scenarios enforce their payoff structures. Let me explain. Telephone games are cooperative in the sense that everyone benefits or loses together. There are many reasons why a given scenario could have this cooperative structure. Familiar examples include the benefits that “I” gain from my kin prospering. A different mechanism is defended by Tomasello (2016) who argues, building on Roberts (2005), for a “stakeholder” view in which agents benefit from the continued success of members of their community who provide them with indirect benefits. Or the payoff structure could be cooperative because agents will withhold future

benefits to those who have provided them with inaccurate information (Mercier and Sperber 2017), or because there’s good reasons for agents to expect other agents to behave in a cooperative manner (Bicchieri 1997, 2006). All of these mechanisms are capable of producing scenarios in which I benefit if and only if you do.

As I noted above, the model I’ve produced doesn’t take a stand on why we find ourselves in a cooperative scenario; the question I’ve asked is simply whether, on finding ourselves in this scenario, we should expect rational agents to be favored or not. It seems plausible that in this particular case, the reason why we find ourselves in a cooperative scenario is that demonstrating our reliability as a testifier is crucial to ensuring the cooperation of future partners. Strictly speaking, however, this aspect of the view hasn’t been tested by my model, which applies regardless as to which mechanisms explain the cooperative structure of the payoffs. I take it that it’s simply a separate question what mechanisms are capable of generating scenarios like the one I’ve described as “stag-hunt-like,” though we need to be careful here as the details of how cooperation is enforced can have an effect on the what strategies the cooperative scenario favors (Bicchieri 2006; Forber and Smead 2015).

6 Conclusion

In this paper, I’ve developed a model of cooperative interaction designed to test whether cooperation can foster epistemic rationality. The results indicate that cooperation can work as advertised by theorists such as Tomasello: at least under some (plausible) conditions, cooperative scenarios favor epistemically rational agents over wishful agents specifically designed to perform as well as them in non-interactive contexts. There’s much more work to be done here—particularly in investigating Tomasello’s specific claims about the role of partner choice and stakeholder value in cooperative scenarios, and the effects that such mechanisms have on the structure of the cooperative games and thus which strategies are favored when—but the results are promising and demonstrate the potential usefulness of game-theoretic tools in investigating the origins and grounds of epistemic rationality.

Appendix: Technical details

Defining agential behavior

Let Pr_s indicate the agent’s subjective probability, Pr_o the objective probability, and u the utility function for that agent. Our rational agents have perfect uptake of the objective probabilities (i.e., for all $s_i \in S$, $Pr_s(s_i|o) = Pr_o(s_i|o)$), act to maximize expected utility, and are “credulous” in the sense that when they receive a report regarding the probability distribution from another agent, their subjective probability assignment is the same as the assignment that they receive (letting Pr_r indicate the probability assignment received: for all $s_i \in S$, $Pr_s(s_i) = Pr_r(s_i)$). Credulity is an idealization, obviously, but it may be a good one (Zollman 2015).

Wishful agents are little more complicated. First, they assign probabilities as follows:

$$Pr_s(s_i|o) = \frac{Pr_o(s_i|o)u(s_i)}{\sum_{s_j \in S} Pr_o(s_j|o)u(s_j)}$$

which is to say that their probability assignment is equivalent to their expected utility for each state normalized so as to behave like probabilities. Similarly, when receiving a probability assignment, they incorporate their utilities into their assignment of probabilities like so:

$$Pr_s(s_i) = \frac{Pr_r(s_i)u(s_i)}{\sum_{s_j \in S} Pr_r(s_j)u(s_j)}$$

In other words, wishful agents are “credulous” in the same sense that rational agents are: they don’t discount the reports of others by taking into account the trustworthiness of the sender or by or incorporating their own priors. Instead, they take the probabilities at “face value,” only incorporating their preferences, just as they do in response to an observation. Finally, in order to avoid double-counting their utilities, they determine actions a little bit differently, namely by picking the action that has the highest “probability” according to their assignment.

Note that this particular way of cashing out what “wishful” means only results in identical behavior with rational agents in non-interactive scenarios due to stipulations about the payoff structure. In particular, it requires that action x and state x are guaranteed to have the same expected utility.

Determining expected utilities

In any given scenario, every agent has an expected utility. So, for example, in a degree-2 cardinality-2 game with one rational agent and one wishful agent mean observation power set to .75, and observation power σ^2 of .05, the expected utility of the rational agent (u_R), wishful (u_W) (wishful), and actor (u_A) $u_R(1, 1) = 1.19$, $u_W(1, 1) = 1.29$, and $u_A(1, 1) = 1.29$, where the numbers in parentheses indicate the number of rational and wishful agents in the chain respectively. To determine these values, I ran one million simulations for each scenario.

The overall expected utility for an agent is then determined by how likely they are to find themselves in each possible scenario, which depends on the agent themselves and the proportion of agents playing the rational strategy. Letting R indicate this latter value, the overall expected utility for a rational agent in any degree-2 cardinality-2 game is

$$u_R = \frac{2}{3} \left[[R * u_R(2, 0)] + [(1 - R) * u_R(1, 1)] \right] + \frac{1}{3} \left[[R^2 * u_A(2, 0)] + [2 * R * (1 - R) * u_A(1, 1)] + [(1 - R)^2 * u_A(0, 2)] \right] \quad (1)$$

Spelling this out: any agent in a cardinality-2 game has a 2 in 3 probability of being a testifier rather than an actor; given that they are a testifier, their probability of being paired with a rational agent is given by R and their probability of being paired with a wishful agent by $1 - R$. These facts allow us to weight the probability of different scenarios that the rational agent can find themselves in, giving us their total expected utility when entering into the game. The expected utility for a wishful agent, while identical in the second term, differs from the above in that the first term is

$$u_W = \frac{2}{3} \left[[R * u_W(1, 1)] + [(1 - R) * u_W(0, 2)] \right] + \dots \quad (2)$$

The difference stems from the features of the individual agent themselves: wishful agents can't find themselves in a group of only rational agents, because they themselves are part of the group (and similarly for rational agents).

The initial model

To run the model itself, I set R to .5 and then allowed it to evolve according to a discrete version of the replicator dynamics. That is, at each step n , $R(n+1)$ was calculated by adding the weighted advantage of rational agents over the average agent like so:

$$R(n+1) = R(n) + R(n)[u_R - [R(n) * u_R + (1 - R(n)) * u_W]] \quad (3)$$

The simulation halted when one of three conditions was met: (a) one strategy was played by 99.99% or more of agents, (b) the percentage of agents playing each strategy after a given round was identical (to the hundredth of a percent) to the percentage playing after the previous round, (c) 1000 rounds had passed. (Why 1000 rounds? I found that there are three different behaviors that I wanted to distinguish: (a) halting at a percentage substantially under 100%, (b) relatively quick but ultimately asymptotic approach of 100%, and (c) very slow movement towards 100%. At 500, it's hard to distinguish between (a) and (c). At 1500 it starts to get hard to distinguish between (b) and (c). 1000 is a happy medium.)

A summary of the results for the first game can be found in table 2.

Complicating the model

The complications introduced in section 4 work as follows. First, I introduced non-cooperators who do not interact with other agents: they simply take a guaranteed payoff, which ensures that their expected utility is equal to $1 + 1 / \text{the cardinality of the game}$ (the second term represents their chance of getting a “preferred” state via this guaranteed payoff). Second, I had cooperators of both wishful and rational variety randomly attempt to cooperate with five other individuals, with the degree of the game depending on how many of those individuals were also cooperators and failure to find any other cooperators resulting in a payoff of 0. This captures the element of risk associated with the cooperative strategy: the more non-cooperators in the population, the higher the probability a cooperator gets nothing. Third, to capture the benefits of cooperation, I increased the payoffs depending on the number of cooperators in the group: it's hard to get five cooperators together—and doing so risks massive information loss—but the rewards for success in such cases are dramatic. Specifically, I set the advantage equal to $1 + \text{the number of total cooperators in the}$

		Cardinality 2									
		Degree 1		Degree 2		Degree 3		Degree 4		Degree 5	
$Pr(s_t o)$		%R	Turns	%R	Turns	%R	Turns	%R	Turns	%R	Turns
.10		16.6	1000	0	45	0	95	0	187	0	511
.20		0.8	1000	0	47	0	107	0	288	0.1	1000
.30		0.7	1000	0	66	0	106	0.1	1000	3.8	1000
.40		100	157	0	49	0	705	0.8	1000	5.5	1000
.50		1.7	1000	0	1000	0.6	1000	8.1	1000	12.4	1000
.60		0.5	1000	40.7	81	58.2	83	67.8	87	73.8	94
.70		100	272	100	176	100	152	100	140	100	134
.80		100	212	100	134	100	115	100	105	100	101
.90		99.1	1000	100	1000	100	1000	100	1000	100	1000

		Cardinality 3									
		Degree 1		Degree 2		Degree 3		Degree 4		Degree 5	
$Pr(s_t o)$		%R	Turns	%R	Turns	%R	Turns	%R	Turns	%R	Turns
.10		2.6	1000	0	84	0	227	0	464	0	1000
.20		87.7	1000	0	73	0	283	0.2	1000	4.9	1000
.30		100	600	0	109	0.1	1000	4.7	1000	8.7	1000
.40		0.6	1000	36.0	97	53.2	94	63.1	124	69.5	129
.50		1.9	1000	74.9	168	86.0	147	90.3	150	92.5	140
.60		100	138	100	88	100	76	100	71	100	68
.70		100	468	100	267	100	223	100	201	100	188
.80		55.2	1000	99.9	1000	99.9	1000	99.9	1000	100	1000
.90		50.0	1	99.8	1000	99.9	1000	99.9	1000	99.9	1000

		Cardinality 4									
		Degree 1		Degree 2		Degree 3		Degree 4		Degree 5	
$Pr(s_t o)$		%R	Turns	%R	Turns	%R	Turns	%R	Turns	%R	Turns
.10		8.4	1000	0	100	0	411	0	1000	1.9	1000
.20		100	437	0	117	0	1000	1.8	1000	6.2	1000
.30		0.7	1000	34.6	130	51.7	120	61.7	167	68.3	185
.40		2.2	1000	80.1	178	89.1	175	92.5	161	94.3	166
.50		100	161	100	100	100	86	100	79	100	75
.60		100	490	100	279	100	232	100	209	100	196
.70		56.4	1000	99.8	1000	99.9	1000	99.9	1000	100	1000
.80		50.0	1	99.6	1000	99.8	1000	99.9	1000	99.9	1000
.90		50.0	1	99.6	1000	99.8	1000	99.9	1000	99.9	1000

Table 2: Summary of simulation results for games of cardinality 2, 3, and 4, where “%R” indicates the percentage of the population playing the rational strategy and “Turns” the number of turns the model took to shut off.

interaction divided by 2. (This particular choice of advantage is largely arbitrary, but the overall pattern of when cooperation is favored doesn't depend on it.) The expected utility of a rational agent is then given by:

$$u_R = [Pr(1 \text{ cooperators}) * \text{ex. payoff in d-1 game} * \text{adv. of d-1 game}] \\ + [Pr(2 \text{ cooperators}) * \text{ex. payoff in d-2 game} * \text{adv. of d-2 game}] + \dots \quad (4)$$

$$= [5(1-N)(N)^4 * u_{R_1} * 1.5] + [10(1-N)^2(N)^3 * u_{R_2} * 2] + \dots \quad (5)$$

Where N is the proportion of the population who are *not* cooperating and u_{R_x} (or, in the wishful case, u_{W_x}) the expected utility of agents in games of degree- x . The extension to the wishful case is straightforward.

Introducing the additional kind of agent required slightly altering the dynamics. Rather than (3), above, $R(n+1)$ and $N(n+1)$ were calculated like so:

$$R(n+1) = R(n) + (1 - N(n)) * R(n) [u_R - [R(n) * u_R + (1 - R(n)) * u_W]] \quad (6)$$

$$N(n+1) = N(n) + N(n) [u_N - [N(n)u_N + (1 - N(n)) * [R(n) * u_R + (1 - R(n)) * u_W]]] \quad (7)$$

The second equation is standard: $N(n+1)$ is calculated by adding the weighted advantage of non-cooperators over the average agent. The first is modified to account for the assumption that the difference between rational and wishful agents is only relevant to cooperators.

A summary of the results for the second game can be found in table 3.

References

- Bicchieri, C. (1997). Learning to Cooperate. In: *The Dynamics of Norms*. Ed. by C. Bicchieri, R. Jeffrey, and B. Skyrms. Cambridge: Cambridge University Press: 17–46.
- (2006). *The Grammar of Society: The Nature and Dynamics of Social Norms*. Cambridge: Cambridge University Press.
- (2017). *Norms in the Wild: How to Diagnose, Measure, and Change Social Norms*. Oxford: Oxford University Press.

		Cardinality 2									
		Initial $N = .1$		Initial $N = .3$		Initial $N = .5$		Initial $N = .7$		Initial $N = .9$	
$Pr(s o)$		%R	%N	%R	%N	%R	%N	%R	%N	%R	%N
.10		0	0	33.8	100	45.1	100	48.8	100	49.9	100
.20		0	0	20.0	100	42.7	100	48.1	100	49.8	100
.30		3.4	0	3.4	0	41.4	100	48.0	100	49.8	100
.40		3.7	0	3.7	0	40.7	100	49.4	100	50.0	100
.50		4.6	0	4.6	0	4.6	0	48.9	100	49.9	100
.60		73.8	0	73.8	0	73.8	0	47.2	100	49.9	100
.70		100	0	100	0	100	0	100	0	50.1	100
.80		100	0	100	0	100	0	100	0	50.1	100
.90		100	0	100	0	100	0	100	0	50.1	100

		Cardinality 3									
		Initial $N = .1$		Initial $N = .3$		Initial $N = .5$		Initial $N = .7$		Initial $N = .9$	
$Pr(s o)$		%R	%N	%R	%N	%R	%N	%R	%N	%R	%N
.10		23.6	100	41.3	100	46.8	100	49.1	100	49.9	100
.20		4.0	0	39.0	100	46.7	100	49.3	100	50.0	100
.30		3.9	0	39.7	100	47.7	100	49.6	100	50.0	100
.40		69.6	0	69.6	0	46.0	100	48.8	100	49.9	100
.50		92.5	0	92.5	0	92.5	0	51.6	100	50.0	100
.60		100	0	100	0	100	0	100	0	50.2	100
.70		100	0	100	0	100	0	100	0	50.1	100
.80		100	0	100	0	100	0	100	0	50.1	100
.90		100	0	100	0	100	0	100	0	50.1	100

		Cardinality 4									
		Initial $N = .1$		Initial $N = .3$		Initial $N = .5$		Initial $N = .7$		Initial $N = .9$	
$Pr(s o)$		%R	%N	%R	%N	%R	%N	%R	%N	%R	%N
.10		32.9	100	43.7	100	47.7	100	49.4	100	50.0	100
.20		34.4	100	45.0	100	48.5	100	49.7	100	50.0	100
.30		54.1	100	48.5	100	48.3	100	49.2	100	49.9	100
.40		94.3	0	94.3	0	94.3	0	94.3	100	50.0	100
.50		100	0	100	0	100	0	100	0	50.1	100
.60		100	0	100	0	100	0	100	0	50.1	100
.70		100	0	100	0	100	0	100	0	50.1	100
.80		100	0	100	0	100	0	100	0	50.1	100
.90		100	0	100	0	100	0	100	0	50.1	100

Table 3: Summary of simulation results for the more complicated games. As before, %R indicates the percentage of rationals; %N indicates the percentage of non-cooperators. The initial proportion of rational to wishful agents in all three tables was 1:1.

- Dogramaci, S. (2012). Reverse Engineering Epistemic Evaluations. *Philosophy and Phenomenological Research* 84.3: 513–30.
- (2015). Communist Conventions for Deductive Reasoning. *Noûs* 49.4: 776–99.
- Douglas, H. (2009). *Science, Policy, and the Value-Free Ideal*. Pittsburgh, PA: Pittsburgh University Press.
- Forber, P. (2010). Confirmation and Explaining How Possible. *Studies in History and Philosophy of Science Part C* 41.1: 32–40.
- Forber, P. and R. Smead (2015). Evolution and the Classification of Social Behavior. *Biology & Philosophy* 30.3: 405–21.
- Gigerenzer, G. (2008). *Rationality for Mortals: How People Cope With Uncertainty*. Oxford: Oxford University Press.
- Graham, P. J. (2020). Assertions, Handicaps, and Social Norms. *Episteme* (online first).
- Hieronymi, P. (2005). The Wrong Kind of Reason. *Journal of Philosophy* 102.9: 437–57.
- Lewis, D. (1969). *Convention*. Cambridge, MA: Harvard University Press.
- McLoone, B. and R. Smead (2014). The Ontogeny and Evolution of Human Collaboration. *Biology & Philosophy* 29: 559–76.
- McLoone, B. et al. (2018). Stochasticity, Selection, and the Evolution of Cooperation in a Two-Level Moran Model of the Snowdrift Game. *Complexity* 2018: 1–14.
- Mercier, H. and D. Sperber (2017). *The Enigma of Reason*. Cambridge, MA: Harvard University Press.
- Morton, J. M. (2017). Reasoning under Scarcity. *Australasian Journal of Philosophy* 95.3: 543–59.
- O’Connor, C. (2020). *Games in the Philosophy of Biology*. Cambridge: Cambridge University Press.
- Okasha, S. (2013). The Evolution of Bayesian Updating. *Philosophy of Science* 80.5: 745–57.
- (2018). *Agents and Goals in Evolution*. Oxford: Oxford University Press.
- Pandit, V., A. Mukhopadhyay, and S. Chakraborty (2018). Weight of Fitness Deviation Governs Strict Physical Chaos in Replicator Dynamics. *Chaos* 28.3: 1–12.
- Ridley, M. (2001). *The Cooperative Gene: How Mendel’s Demon Explains the Evolution of Complex Beings*. New York: Free Press.
- Roberts, G. (2005). Cooperation Through Interdependence. *Animal Behavior* 70.4: 901–8.

- Sharadin, N. (2018). Epistemic Instrumentalism and the Reason to Believe in Accord With the Evidence. *Synthese* 195: 3791–809.
- Skyrms, B. (1996). *Evolution of the Social Contract*. Cambridge: Cambridge University Press.
- (2004). *The Stag Hunt and the Evolution of Social Structure*. Cambridge: Cambridge University Press.
- (2010). *Signals: Evolution, Learning, and Information*. Oxford: Oxford University Press.
- Sterelny, K. (2012). *The Evolved Apprentice: How Evolution Made Humans Unique*. Cambridge, MA: The MIT Press.
- Tomasello, M. (2014). *A Natural History of Human Thinking*. Cambridge, MA: Harvard University Press.
- (2016). *A Natural History of Human Morality*. Cambridge, MA: Harvard University Press.
- (2017). *Becoming Human: A Theory of Ontogeny*. Cambridge, MA: Harvard University Press.
- (2020). The Ontogenetic Foundations of Epistemic Norms. *Episteme* (online first).
- Tomasello, M. et al. (2012). Two Key Steps in the Evolution of Cooperation: The Interdependence Hypothesis. *Current Anthropology* 53.6: 673–92.
- Zollman, K. J. S. (2015). Modeling the Social Consequences of Testimonial Norms. *Philosophical Studies* 172.9: 2371–2383.
- (2017). Learning to Collaborate. In: *Scientific Collaboration and Collective Knowledge: New Essays*. Ed. by T. Boyer-Kassem, C. Mayo-Wilson, and M. Weisberg. Oxford: Oxford University Press: 65–77.