

Social deliberation: Nash, Bayes and the partial vindication of Gabriele Tarde*

J. McKenzie Alexander

Department of Philosophy, Logic and Scientific Method
London School of Economics and Political Science

Abstract

At the very end of the 19th century, Gabriele Tarde wrote that all society was a product of imitation and innovation. This view regarding the development of society has, to a large extent, fallen out of favour, and especially so in those areas where the rational actor model looms large. I argue that this is unfortunate, as models of imitative learning, in some cases, agree better with what people actually do than more sophisticated models of learning. In this paper, I contrast the behaviour of imitative learning with two more sophisticated learning rules (one based on Bayesian updating, the other based on the Nash-Brown-von Neumann dynamics) in the context of social deliberation problems. I show for two social deliberation problems, the Centipede game and a simple Lewis sender-receiver game, that imitative learning provides better agreement with what people actually do, thus partially vindicating Tarde.

By the end of the 19th century, hopes that we were getting close to a complete scientific understanding of the world were running high in some quarters. Perhaps the greatest indicator of such hopes was when Lord Kelvin declared in his address to the British Association for the Advancement of Science, “There is nothing new to be discovered in physics now. All that remains is more and more precise measurement.” And these hopes regarding our ability to explain were not just limited to the physical world, but extended to the social world, the world of human behavior, as well. For in 1890, just ten years prior to Lord Kelvin’s pronouncement of the end of physics, Gabriele Tarde, a French sociologist, stated in similar grandiloquent fashion that he had identified the fundamental forces governing society: “What is society? I have answered: society is

*Published in *Episteme: A Journal of Social Epistemology*, vol. 6, no. 2, pp. 164–184 (2009).

imitation.” (Tarde, 1903, pg. 74) In a later expansion of the same idea, we find: “[a]ll resemblances of social origin in society are the direct or indirect fruit of the various forms of imitation.”¹

Of course, from a contemporary perspective most declarations exhumed from the dustbin of history appear ridiculous. (Everything is made of water?) But whereas we might be willing to overlook the hubris of Lord Kelvin, given that his remarks about the “end of physics” must be viewed alongside his prescient identification of the two outstanding issues which led to modern physics², there seems less reason to exercise a similar degree of charity for M. Tarde. Society is not just a product of imitation and innovation by individuals. How could anyone think it was? Our social lives are shaped by a rich variety of factors which Tarde’s theory omits, such as social norms, strategic reasoning, empathic concern for the well-being of others, the requirements of duty, obligation, and so on. These factors shape our behaviour and thought in ways which cannot be reduced to mere imitation.

Yet even though Tarde’s theory, as a complete account of society, is false, that does not mean there are no insights worth preserving. Indeed, in recent years there has been a small resurgence of interest³ in the writings of this French sociologist who was, at the time of his death, considered one of the greats among Comte, Darwin and Spencer (Millet, 1970). In what follows, I offer an attempt to partially rehabilitate, and partially vindicate, the views of Tarde regarding the importance of imitation for society. This partial vindication occurs in a very limited sense: I will argue that imitation, when used as a heuristic by boundedly rational individuals, selects the socially optimal outcome in several social deliberation problems much more readily than two other types of deliberative procedures. (The other two types of deliberative procedures being derived from the work of Nash and Bayes, hence the title of this paper.)

What this means is that imitation can be seen as a method for generating and supporting some socially beneficial practices, even if it is not the universal

¹In these passages, Tarde speaks as if imitation were the only social force in operation, but that is obviously insufficient: imitation requires something to imitate, and so there must be a second force at work which generates innovative behaviours. Tarde was well aware of this, even if he omitted the role of invention at times. He does state rather earlier that “[s]ocially, everything is either invention or imitation” (Tarde, 1903, pg. 3). I suspect his tendency to emphasize imitation over invention was due to his belief that imitative behavior was subject to law, whereas invention was not. See, for example, *The Laws of Imitation*, pg. 142.

²In a lecture entitled “Nineteenth-Century Clouds over the Dynamical Theory of Heat and Light,” Lord Kelvin noted that the current physics of the time could not provide a satisfactory account of black body radiation and the Michelson-Morley experiment. These two outstanding problems eventually led to the development of quantum mechanics and relativity theory.

³Latour (2002), for example, argues that Tarde can be viewed as a intellectual precursor to his “actor-network” theory.

social *explanans* that Tarde thought. That might not sound all that interesting or important, but it becomes so once we factor in *when* people choose to imitate. In the models I present, imitation occurs when individuals, who are purely motivated by the desire to maximize their personal gain, believe they can do better by adopting the behaviour of another. Contrast this with the well-known fact that other decision rules which seek to maximize individual gain often fail to generate socially optimal outcomes, as in the tragedy of the commons, the prisoner’s dilemma, and the centipede game. Imitation, then, provides a way for individuals to strive to maximize their own personal gain in a way which does not preclude arriving at socially optimal outcomes. That, I think, is a claim worth noting.

1. Social deliberation problems.

In what follows, I shall take a *social deliberation problem* to refer to problems of the following form: a population of agents faces a multitude of interdependent, noncooperative, two-person decision problems, with the special property that each individual, when he chooses an action, has to use that action regardless of whom he interacts with.⁴ It is a *social* deliberation problem in the sense that I cannot choose to condition my behaviour on the identity of persons I am going to interact with. (Think of it as my having to choose a single “face” which I present to all members of society.) It is a social *deliberation* problem in the sense that each individual does not have a fixed idea in her mind as to what she is going to do, and so will modify her beliefs in response to others.

Let us formulate this more precisely. Let $P = \{1, \dots, N\}$ denote the population of agents and let M be the payoff matrix for the two-player decision problem as follows:

$$M = \begin{pmatrix} \langle r_{11}, c_{11} \rangle & \cdots & \langle r_{1n}, c_{1n} \rangle \\ \vdots & \ddots & \vdots \\ \langle r_{n1}, c_{n1} \rangle & \cdots & \langle r_{nn}, c_{nn} \rangle \end{pmatrix}.$$

Note that the payoffs for the two players may not be equal, but it is required

⁴One might wonder whether the requirement that the decision problem be noncooperative unduly restricts the kinds of problems that can be treated as social deliberation problems. In principle I don’t see why it needs to, if one is willing to adopt the approach of the Nash program to embed cooperative game theory within noncooperative game theory. Another concern may be with why I restrict attention to two-player games. Essentially, it makes the formal models simpler; the requirement that players use a single strategy with everyone they interact with effectively transforms the “real” game from a two-player game to one where people “play the field” (in a certain sense).

that the same number of strategies be available for both, regardless of whether they play as row or column.

One important feature about society is that it has *structure*. Social structure can be modeled in a variety of ways, but the most important aspect of it is that it endures and constrains individual choice and action. Adopting a model of social structure which I have used elsewhere (Alexander, 2007), let us model the structure of society by a directed graph $G = \langle P, E \rangle$, where the set of directed edges E represents a binary relation of some particular social importance (such as X being a friend of Y , X being an acquaintance of Y , and so on). Given a particular individual i , the set of players with whom i shares an edge with are the *neighbours* of i . The neighbours of a player are the individuals with whom he plays the game. If i and j are connected by an edge pointing from i to j , that means when the two play a game, i plays as Row and j plays as Column.

People deliberate when they are uncertain about what to do. The state of uncertainty of each player is represented as a probability distribution over the possible actions available to him. (I shall treat “action” as synonymous with “strategy” given that we are working in a game theoretic context.) We can think of the state of uncertainty of player i , at time t , as a vector $\vec{p}_i(t) = \langle p_{i_1}(t), \dots, p_{i_n}(t) \rangle$, where $p_{i_n}(t)$ denotes the probability that i assigns to action n at time t . This probability may be interpreted as a measure of how “desirable” that action appears to i at the time.

How do people deliberate, and how do people end up revising their state of uncertainty as a consequence of their deliberations? It is easy to imagine a variety of ways this might happen. It turns out that there is a natural way to extend the two-person deliberative dynamics of Skyrms (1990) to the socially structured setting envisioned here.⁵ Let us assume that people’s deliberation over how to revise their state of uncertainty occurs in two stages: first, a player deliberates about how she would revise her state of uncertainty for each pairwise interaction with a neighbour, given what she knows about him or her. Second, once the player has determined what each of these pairwise revisions would be, she then proceeds to aggregate, or pool, these multiple revisions into a single state of belief.

Again, let us state this more precisely. Suppose that $\eta_i = \{i_1, \dots, i_j\}$ denotes the set of neighbours of player i . If $\vec{p}_i(t)$ denotes the state of uncertainty of player i at time t , let $\vec{p}_{i,i_k}(t+1)$ denote the state of uncertainty that player i *would* have at time $t+1$, if she revised her current state of uncertainty given just what she knows about i_k . (Hence $\vec{p}_{i,i_k}(t+1)$ also denotes what i ’s future state

⁵Another version of this deliberative model can be found in Alexander (2009). However, the primary aim of that paper concerned the outcomes of social deliberation in simple coordination games like the Driving game, Battle of the Sexes, and Chicken.

of uncertainty would be if she had only one neighbour, namely i_k .)

Once player i has calculated the pairwise refinements $\vec{p}_{i,i_k}(t+1)$ for all of her neighbours $i_k \in \eta_i$, she then aggregates these possible refinements into a single probability distribution over action. The aggregation rule I assume players use is the following⁶:

$$\vec{p}_i(t+1) = \frac{1}{j} \sum_{k=1}^j \vec{p}_{i,i_k}(t+1).$$

This is a linear pooling method for aggregating probabilities; it also assigns equal weights to each of the possible pairwise refinements of i 's current state of uncertainty. Using equal weights makes sense if all of i 's neighbours are equally important to her. (One could easily generalize this by attaching weights to edges to indicate how important player j was to player i .)

Assuming that players use a linear pooling method for aggregating their possible future states of uncertainty makes sense because such methods are the only ones which satisfy the following requirements (Lehrer and Wagner, 1981):

1. The aggregate probability player i assigns to strategy σ_k in his state of uncertainty at time $t+1$ depends *only* upon the probability i assigns to σ_k in each pairwise refinement to his state of uncertainty.
2. If a player assigns probability zero to a strategy in each pairwise refinement to his state of uncertainty, then that player assigns probability zero to that strategy in his aggregate state of uncertainty.

These are reasonable requirements to impose.

What deliberative rule do people use when calculating the pairwise refinements of their state of uncertainty? I shall consider two: the first employs what is known as the Nash-Brown-von Neumann dynamics, as it derives from the function Nash used in his proof of his fixed point theorem. (I shall refer to this as just the "Nash dynamics" for simplicity.) The second rule is a variant of Bayesian updating.⁷ Both dynamical rules provide an approximation of how rational agents would deliberate over what to do. And both of these dynamical rules are in keeping with what a view of human agents as sophisticated deliberative agents would endorse.

On the other hand, Gabriele Tarde thought that the deliberative dynamics underlying society were rather different from the above. He thought that society existed as a consequence of *imitation*. We can easily formulate a social network

⁶Note that the value of j depends on the number of neighbours player i has. I have suppressed this dependence to make the notation more clear.

⁷Definitions of these two rules can be found in appendix A.

model of imitative learning. As before, assume that there is a social network which determines who interacts with whom. Let us assume that each person plays the game with every one of his neighbours, receiving a score equalling the sum of the individual payoffs from each pairwise interaction. At the end of each round of deliberation, each player i looks at his set of neighbours and adopts the strategy used by his neighbour who did the best (assuming, of course, that this payoff exceeds the payoff of player i). This dynamic is known as “Imitate the Best” and has been suggested as a useful heuristic for boundedly rational individuals (see [Gigerenzer and Selten, 2001](#)).

These two different models of social deliberation impose substantially different requirements on what people know.⁸ The Nash dynamics and Bayesian dynamics assume that, when i and j are connected by an edge, that each player’s full state of uncertainty is common knowledge. They also assume that the deliberative rule used by i and j is common knowledge. On the other hand, Imitate-the-Best makes no such assumptions. Imitate-the-Best does not assume that each player’s state of uncertainty is known by anyone else in the population. Moreover, when a person adopts a new strategy through imitation, the new strategy is, in this model, necessarily a pure strategy rather than a probability distribution. Why? If i imitates j , player i adopts the last move made by player j in the game. But the last move of player j is a pure strategy.⁹

We have, then, two different kinds of models of social deliberation. One model treats individuals as highly rational, with considerable amounts of common knowledge about their neighbours, willing to use sophisticated aggregation techniques to try to find the optimal outcome. The other model treats individuals as boundedly rational, with very little knowledge about their neighbours, who simply imitate the best. In the next two sections I show that imitative learning, rather than the more sophisticated models of social deliberation, is better suited for producing socially optimal outcomes in the Centipede game and in a Sender-Receiver game.

2. The Centipede game.

The Centipede game (see [Rosenthal, 1981](#)) is a well-known example of an interpersonal decision problem in which the traditional game theoretic analysis

⁸There are only two different models because the knowledge assumptions for the Nash dynamics and the Bayesian dynamics are the same.

⁹Think of it this way: if an agent only sees a finite sequence of actual actions you have made, that agent cannot reconstruct what your underlying probability distribution is, since any finite sequence is compatible with any probability distribution. If the agent were to try to track the frequency with which you have chosen certain actions, we do not have a model of *imitation* but rather *inductive learning*.

conflicts with what our intuitions suggest as the way to play. Figure 1 illustrates a six-stage Centipede game.¹⁰ Player I begins at the root node, located at the far left, and has two choices: either *take* the amount available, or *pass* to the other player. If player I chooses to pass, Player II faces the exact same choice: take what is available, or pass control back to player I. Inspection of the payoffs show that the socially optimal outcome (here, a collective payoff of 14) occurs when both players always choose Pass. However, if one solves the game using backwards induction, it turns out that what player I should do is choose *Take* on the very first move, giving himself a payoff of 2 and player II a payoff of zero.¹¹




Figure 1: A six-stage Centipede game.

Although backwards induction recommends that player I take on the first move, this conflicts with the intuitions of some that even rational players should move to the right in the beginning, at least for a while, before choosing to move down. These intuitions are borne out by experiment. McKelvey and Palfrey (1992) report that, in a six-stage centipede game, only 1% of the players chooses *Take* on the first move. When the game reaches the final stage, 15% of the time the last player chooses *Pass*, thereby playing a dominated strategy but, at the same time, producing the socially optimal outcome.¹²

¹⁰The game takes its name from the fact that in the original formulation there were a hundred such segments. The same tension exists in the shorter version.

¹¹Consider the last choice node for player II. If she chooses *pass*, she receives a payoff of 6 but if she chooses *take*, she receives a payoff of 7. A rational agent interested in maximising her personal gain will choose *take* (thus giving player I a payoff of 5). Player I knows this, and so at his last choice node will prefer to preempt player II's decision by choosing *take*, since that gives him a payoff of 6, which is greater than 5. Continuing this reasoning leads to the outcome that player I will choose *take* at the very start of the game.

¹²Somewhat curiously, people continue to deviate from the game theoretic prediction even in *constant-sum* centipede games. (A constant-sum centipede game is one where the pot, instead of growing with each *Pass* as in figure 1, remains constant over time; in these games, choosing *Pass* repeatedly has the effect of increasing the amount given to player II.) In a later paper, Fey et al. (1996) report on a number of constant-sum experiments they conducted at CalTech, Pasadena City College, and the University of Iowa. In these experiments, people chose *Take* as their first move approximately 59% of the time. This is not that surprising, given that the resulting payoffs when player I chose *Take* led to an equal share of the pot (payoffs of 1.60 each). But notice what

Experimental results such as these have been viewed as showing that there is an important mismatch between the outcomes of the traditional analysis and what people both do and think they ought to do. Martin Hollis offers the following trenchant critique:

“The Centipede seems to me to force the basic issue neatly. One could still shrug one’s shoulders and comment that, since the logic clearly tells the first player to open by playing down, there is no more to be said. But, whereas a similar shrugging off of a mutually inferior outcome in the Prisoner’s Dilemma might be a fair comment on a dismal fact of real life, the Centipede is a scandal for Game Theory.” (Hollis, 1994)

What’s his proposed solution to the scandal? Hollis thinks that it requires replacing the underlying model of human agent, swapping *homo economicus* for the more socially sensitive, norm-based, rule-concerned, other-focused *homo sociologicus*.

Saying that the Centipede game is a scandal for game theory strikes me as a bit hyperbolic, but it does raise the question of to what extent we can reconcile observed behavior with the maximizing assumptions underlying game theory. Perhaps moving to an evolutionary game theoretic perspective may help. Let us now compare the outcomes of the various models of social deliberation introduced so far.

Figure 2 illustrates the outcome of a process of social deliberation for a group of 13 agents who play a ten-stage Centipede game. The social structure used is a ring with the direction of edges selected so as to ensure each agent plays the game once in the role of player I and once in the role of player II. The state of uncertainty of an agent is represented using a pie chart, with the size of the i th wedge reflecting the probability assigned to action i by the agent. In the simulation shown, everyone in the population initially is disposed to put probability 1 on choosing *Take* on stage 10.¹³ Under the Nash dynamics, social deliberators immediately begin adjusting their beliefs so as to move away from the socially optimal strategy, converging to the traditional game-theoretic outcome.

Figure 3 illustrates the process by which this happens from the point of view of one of the agents. (To make the following discussion more clear, let us denote

this implies: 41% of the time, player I elected to *Pass*, apparently favouring an *unequal* allocation of payoffs giving more to player II than himself.

¹³There is good reason for choosing this as the initial state of the population. The Nash dynamics have the property that a player who initially assigns zero probability to an action may, after revision, assign positive probability to that action. (This is an important point of difference between the Nash dynamics and Bayesian updating.) If a population of agents who initially assign probability 1 to moving to the far right of the Centipede game evolve to another outcome, this shows that the socially optimal outcome will never evolve.

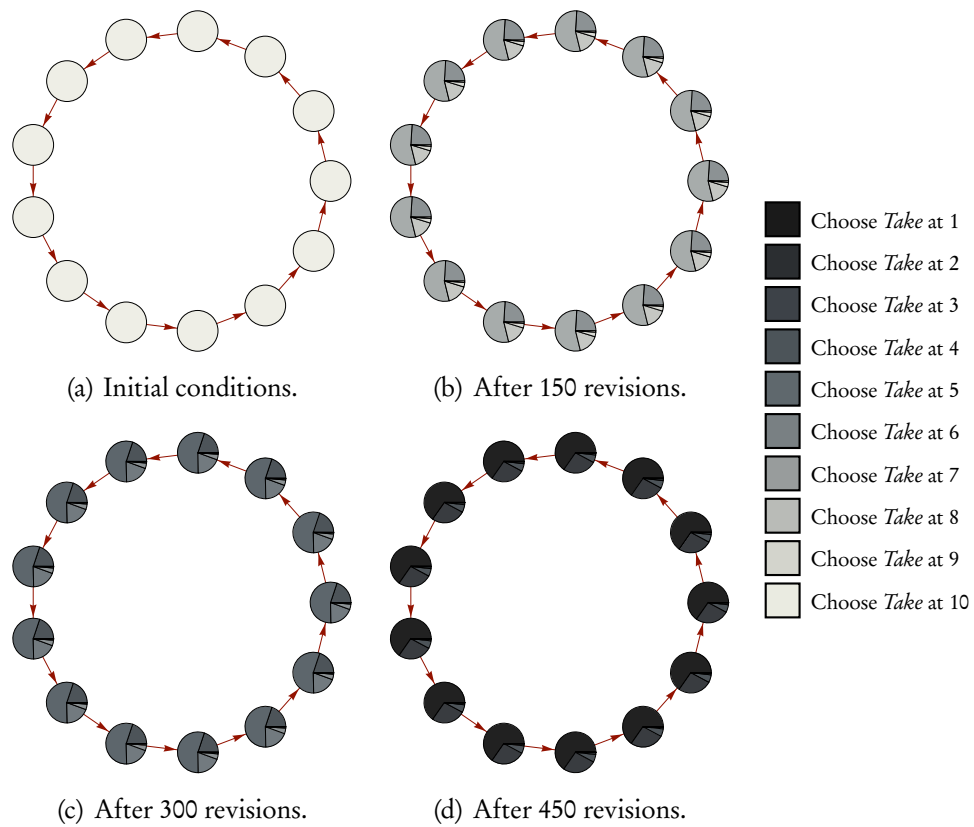


Figure 2: The evolution of the traditional game-theoretic outcome for the Centipede Game under the Nash dynamics (with an index of caution of 10).

the action *Choose Take at Stage n* by S_n .) Suppose that I and my neighbours initially begin by assigning probability one to S_{10} . When I revise my beliefs under the Nash dynamics, the first thing I need to do is calculate the covetability of each possible action available to me. When my neighbours assign probability one to S_{10} , the only action with positive covetability is S_9 . (By symmetry, my neighbours will conclude the same thing.) This means that the new state adopted by both I and my neighbours will be one which puts some small probability ϵ on the action S_9 and probability $1 - \epsilon$ to the action S_{10} .

The next time I revise my beliefs, I compute that both S_8 and S_9 have positive covetability, and so I adopt a new state which assigns positive probability to the actions S_8 , S_9 , and S_{10} . Again, by symmetry my neighbours will do the same thing. As figure 3 illustrates, these three actions will be the only ones to which I assign positive probability for some time; each successive iteration of the social deliberation process will result in me reducing the probability I assign to S_{10} and increasing the probability I assign to S_8 and S_9 . Around the fiftieth revision, though, I find the action S_9 ceases to have a positive covetability, and at this point I will start transferring probability away from both S_9 and S_{10} to S_8 . Eventually it will be the case that the action S_7 will have positive covetability for the first time, and at this point I start increasing the probability of both S_7 and S_8 at the expense of S_9 and S_{10} . (Inspection of figure 3 reveals that this occurs around the 70th iteration of the social deliberation process.) In this fashion, I eventually work my way to putting probability one on S_1 .¹⁴

The general tendency for Nash deliberators to move towards the game-theoretic solution is not affected by variations in social structure. Figure 4 illustrates the outcome of a social deliberation process on a randomly structured social network. Notice that the variability of the number of neighbours does influence the particular probability distribution agents adopt as they revise their beliefs, but it does not affect the long-term result of moving towards the action *Choose Take at Stage 1*. The overall moral of the story is clear: the deliberative outcomes generated by the Nash dynamics do not correspond with what people actually do when faced with the Centipede game.

What happens if people are Bayesian deliberators, rather than Nash deliberators? It turns out that the overall qualitative result is basically the same: Bayesian deliberation leads the population towards the traditional game-theoretic outcome in the Centipede game.¹⁵ Both models of deliberation prove to be equally

¹⁴Convergence to putting probability 1 on S_1 only occurs in the limit.

¹⁵With the proviso that the initial belief state of every agent assigns some positive probability to every possible action. Whereas the Nash dynamics can cause an agent to assign nonzero probability to an action that was initially assigned zero probability, any action assigned zero probability will, under the Bayesian dynamics, always be assigned zero probability in the future. The qualitative result holds in the following sense: suppose that the population starts out




Figure 3: A time-series plot illustrating the evolution of the probability distribution for one of the Nash deliberators in figure 2. The x -axis labels indicate which stage in the deliberative process the bar represents, and each bar represents the probability distribution held by the individual at a given time. The colour-coding is the same as in figure 2.

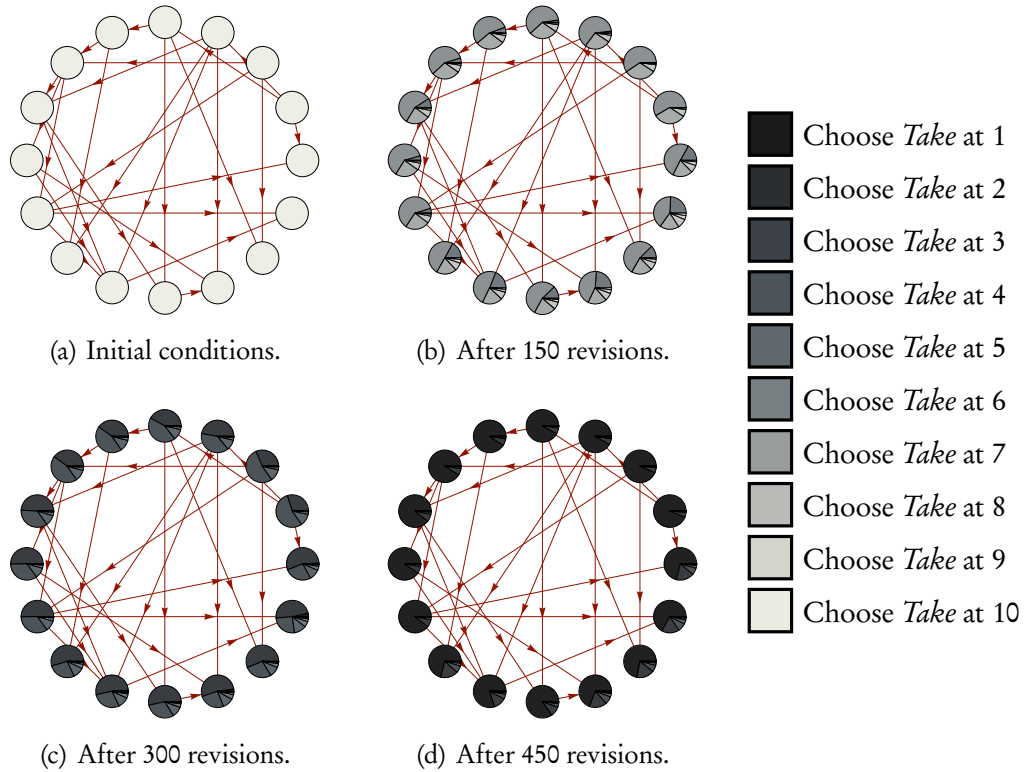


Figure 4: The evolution of the traditional game-theoretic outcome under the Nash dynamics (with an index of caution of 10) on a random network.

poor predictors of what people actually do.

When the social deliberation process takes place using imitative learning, the story is radically different. Figure 5 illustrates the outcome of one simulation using imitative learning where the underlying social network is a lattice.¹⁶ Whereas the previous two deliberative methods proved so hostile to the socially optimal outcome that a population which started in the socially optimal state would quickly leave it, here the situation is entirely the reverse. If the population begins in the state where everyone follows the strategy *Choose Take at Stage 1* and people experiment with new strategies, the population will quickly *leave* the socially inefficient state and move towards the socially optimal one. In the simulation of figure 5, new strategies are introduced with a probability of 2.5% and within 75 iterations virtually everyone in the population has adopted *Choose*

with every agent assigning probability $1 - \epsilon$ to S_{10} and probability $\frac{1}{9\epsilon}$ to S_1, \dots, S_9 . Bayesian deliberators will modify their beliefs so as to move towards the game-theoretic solution.

¹⁶Individuals interact with their eight nearest neighbours and learn from the same group via imitation.

*Take at Stage 10.*¹⁷



Figure 5: Imitative learning leads to the adoption of the socially optimal outcome in the Centipede game.

How likely is it that imitative learners will adopt the socially optimal outcome in the Centipede game? One way of answering this question is through simulation. Because imitative learning is sensitive to the shape of the social network, we will need to sample a variety of different network topologies to control for this dependence.¹⁸ Figure 6 lists the results from 1,000 simulations on random connected networks of 150 agents.¹⁹ Each simulation was run for

¹⁷That is, each individual in the population has a 2.5% chance of replacing his current strategy with a randomly chosen one. With a population of 10,000, approximately 250 individuals will experiment with a new strategy every iteration. Less frequent experimentation rates will still lead to convergence to the socially optimal state (provided that the social network is a lattice); the only substantive difference will be how long it takes convergence to occur.

¹⁸To see why imitative learning depends on the shape of the network, recall the definition of imitate-the-best: a player P adopts the strategy of the person in their neighbourhood who received the highest payoff (provided this payoff was greater than P 's payoff). In an irregular social network, some agents will have more neighbours than others. Agents with more neighbours are more likely to have their strategy adopted by others simply because they engage in more interactions.

¹⁹The random networks were generated with a 3% edge probability with the direction determined by a coin flip. That is, each possible edge had a 3% of being selected for inclusion; if it

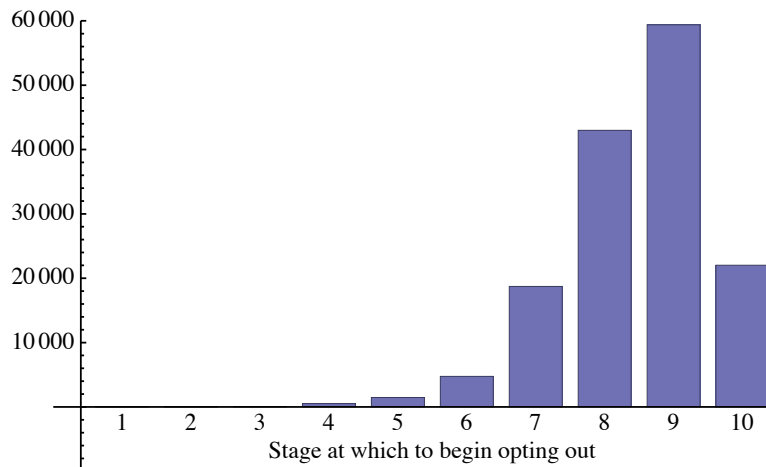


Figure 6: Aggregate results for imitative learning over 1,000 simulations.

300 iterations until convergence (or near-convergence) occurred and the surviving strategies counted. The height of each bar in figure 6 indicates how many agents, out of all 1,000 simulations, followed that strategy after 300 iterations. There are several points worth noting. First, virtually no agent chose *Take* before stage 4. Second, although the strategy which produces the socially optimal outcome appears quite often among the surviving strategies, it is by no means the only such strategy. While the vast majority of agents would choose *Pass* for the first five stages of the Centipede game, once the game entered the sixth stage, agents become increasingly likely to choose *Take*.

The true test of any model is how well it accounts for the experimental data. How does modelling people as boundedly rational imitative learners fare? Recall that McKelvey and Palfrey (1992) found that approximately 15% of subjects in experiment would choose the socially optimal outcome in the Centipede game. The socially optimal outcome in these models occurs when agents choose *Take* in stage 10.²⁰ According to figure 6, approximately 20,000 individuals over the 1,000 simulations followed this strategy at the end of the simulation. Since the total number of strategies counted was 150,000, the socially optimal outcome occurred $\frac{20,000}{150,000} = 13.3\%$ of the time. Perhaps we do not need to replace *homo economicus* with *homo sociologicus* in order to account for what people do in the

was selected, a coin flip determined whether the edge went from *A* to *B* or from *B* to *A*. Since there are $150 \times 149 = 22,350$ possible edges on a graph containing 150 nodes, each graph contained approximately 671 edges. Once a graph was randomly generated, it was tested to ensure it was connected. If it was not connected, that graph was thrown away and another random graph generated. Initial strategies were randomly assigned to players from a randomly chosen distribution.

²⁰Choosing *Pass* at stage 10 was not a option.

Centipede game. It may be that we just need to find the right model of how people try to maximise.²¹

3. Lewis Sender-Receiver games.

Another interesting family of social deliberation problems to consider are the two-player sender-receiver games introduced by David Lewis in *Convention* as a model for the emergence of language.²² In a sender-receiver game, Nature chooses a state of the world and reveals it to one player, known as the Sender, who then sends a signal to a second player, known as the Receiver. Upon receipt of the signal, the Receiver performs an action. If the action is appropriate given the state of the world, both the Sender and Receiver get a payoff of one; if the action is inappropriate, the payoff is zero. Figure 7 shows the extensive-form for a sender-receiver game with two states of the world, two signals, and two actions.




Figure 7: A Lewis sender-receiver game with two states of the world, two signals, and two actions.

If agents can play the sender-receiver game as both Sender or Receiver, then the strategy they use must specify how they will act in either role. A sender-receiver game with two states of the world, two signals, and two actions has

²¹It may even be the case that we need not worry about the specific method people use to maximise. Smead (2008) uses a Moran process to model behaviour in the Centipede game and finds that it, too, gives rise to populations in which people choose *Pass* in the first few stages of the game.

²² N -player versions of sender-receiver games exist as well. See Skyrms (2008) for a nice introduction to these games, along with a discussion of some of the peculiarities that arise.

sixteen possible strategies.²³ Of these, only two have the property that they are *signalling systems* according to Lewis's definition.²⁴ The two Lewis signalling systems are:

1. Send σ_1 in state S_1 , and σ_2 in S_2 ; do A_1 upon receipt of σ_1 , and A_2 upon receipt of σ_2 .
2. Send σ_2 in state S_1 , and σ_1 in S_2 ; do A_1 upon receipt of σ_2 , and A_2 upon receipt of σ_1 .

There are a number of ways one could implement the Sender-Receiver game in a model of social deliberation. Nature could select a randomly chosen state of the world for all pairwise interactions in a given round. Alternatively, it could be the case that, for each pairwise interaction, Nature selects a randomly chosen state of the world. Or we could even consider how individuals would revise their beliefs using the *expected* outcomes of their interactions with their neighbour. In the following, I assume that people revise their beliefs using the expected value of their interaction with their neighbour.

What happens when a population of Nash (or Bayesian) deliberators plays the Sender-Receiver game? Let us consider, as before, the simplest possible social network: k Nash (or Bayesian) deliberators situated on a ring, with the direction of the edges such that each person plays the game once as Sender and once as Receiver.²⁵ As one might expect, quite often the population converges to one of the two Lewisian signalling systems. But this does not always happen. If the ring happens to have an even number of players (say six, for sake of argument), it can happen that players 1, 3, and 5 may converge to one strategy in the signalling game with players 2, 4, and 6 adopting another. How can this be rational?

In an environment where interactions are constrained by a social network, the Lewis signalling systems are not the only signalling systems which allow communication. Consider, for example, the following pair of strategies:

Me. When Sender, use: $(S_1, S_2) \mapsto (\sigma_1, \sigma_2)$.
When Receiver, use: $(\sigma_1, \sigma_2) \mapsto (A_2, A_1)$.

²³There are four possible strategies to use as Sender: $(S_1, S_2) \mapsto (\sigma_1, \sigma_1)$, $(S_1, S_2) \mapsto (\sigma_1, \sigma_2)$, $(S_1, S_2) \mapsto (\sigma_2, \sigma_1)$, or $(S_1, S_2) \mapsto (\sigma_2, \sigma_2)$. Likewise, there are four possible strategies to use as Receiver: $(\sigma_1, \sigma_2) \mapsto (A_1, A_1)$, $(\sigma_1, \sigma_2) \mapsto (A_1, A_2)$, $(\sigma_1, \sigma_2) \mapsto (A_2, A_1)$, or $(\sigma_1, \sigma_2) \mapsto (A_2, A_2)$. Any Sender strategy may be paired with any Receiver strategy, given sixteen possible strategies for the game.

²⁴Lewis's definition of a signalling system is somewhat restrictive in that it excludes combinations of strategies which may nevertheless be perfectly successful at communicating. We shall see that both the Nash and Bayesian dynamics may converge to these nonstandard signalling systems.

²⁵As before I only consider the case where the entire population is composed of either Nash deliberators or Bayesian deliberators; I do not consider heterogenous populations.

You. When Sender, use: $(S_1, S_2) \mapsto (\sigma_2, \sigma_1)$.
 When Receiver, use: $(\sigma_1, \sigma_2) \mapsto (A_1, A_2)$.

Suppose that I am the Sender. In state S_1 I send signal σ_1 , and you respond to that signal by performing A_1 . Likewise, in state S_2 I send signal σ_2 , and you perform action A_2 . The first component of my strategy allows me to signal successfully with the second component of yours. Now suppose that you are the Sender. In state S_1 you send signal σ_2 , and I respond to that signal by performing A_1 . Likewise, in state S_2 you send signal σ_1 , and I respond by performing action A_2 . The first component of your strategy allows you to signal successfully with the second component of mine. Yet the convention used when I am the Sender is the opposite of the convention used when you are the Sender! This is the signalling analogue of you speaking to me in German (and I listen in German), but I speak in Russian (and you listen in Russian). When interactions are constrained via a social network, additional signalling systems exist besides those identified by Lewis.

There are three odd properties about these new signalling systems, though. First, they require that the strategies be very carefully distributed across the population. Second, not every social network allows these new signalling systems to be used if the socially optimal state of affairs is to be achieved. (There is no way that the abovementioned strategies could be used on a ring containing seven players, for example, while at the same time allowing agents to always communicate successfully.) Third, these new signalling systems cannot communicate with members of their own kind. The Lewis signalling systems, on the other hand, can.

Other odd deliberational outcomes exist with both the Nash and Bayesian dynamics. Figure 8 illustrates one such outcome for the Nash dynamics after the players have been deliberating for 20,000 iterations. Here, players 1, 2, and 3 have effectively converged to one of the nonstandard signalling systems discussed earlier.²⁶ Players 6 through 11 have adopted one of the Lewisian signalling systems. Consider, though, the probability distributions adopted by players 4 and 5, and 12 and 13. These four individuals exist between two regions of players who can communicate perfectly with each other. However, their probability

²⁶By “effectively converged,” I mean that they have all assigned more than 99.5% probability to one strategy. In response to those who are inclined to object that these results are uninformative since it may be the case that all individuals in the population converge to a pure Lewis signalling system in the limit, I offer two simple remarks. First, that in the long run we are all dead. Results establishing that a certain outcome holds in the limit are only of practical interest if significant progress towards the limiting outcome can occur within a reasonable amount of time. (A similar point was argued for in [Vanderschraaf and Alexander \(2005\)](#).) Second, given this, the fact that socially inefficient states can persist for 20,000 iterations of the Nash dynamics surely counts as a blow against the Nash dynamics as a mechanism for belief revision.




Figure 8: The Nash dynamics may generate socially suboptimal outcomes in the Sender-Receiver game (state after 20,000 iterations, index of caution of 25).

distributions include strategies which fail to differentiate between states of the world, and hence are ineffective for communicating.

Consider agent 4, for example. His state of belief allocates probabilities over strategies as follows²⁷:

Probability	Strategy	
	Sender	Receiver
0.576843	$\{S_1 \rightarrow \sigma_1, S_2 \rightarrow \sigma_2\}$	$\{\sigma_1 \rightarrow A_2, \sigma_2 \rightarrow A_1\}$
0.293552	$\{S_1 \rightarrow \sigma_1, S_2 \rightarrow \sigma_2\}$	$\{\sigma_1 \rightarrow A_1, \sigma_2 \rightarrow A_2\}$
0.0648024	$\{S_1 \rightarrow \sigma_1, S_2 \rightarrow \sigma_2\}$	$\{\sigma_1 \rightarrow A_1, \sigma_2 \rightarrow A_1\}$
0.0648024	$\{S_1 \rightarrow \sigma_1, S_2 \rightarrow \sigma_2\}$	$\{\sigma_1 \rightarrow A_2, \sigma_2 \rightarrow A_2\}$

The strategy receiving the greatest weight is one of the nonstandard signalling systems, and is the one compatible with player 3. The strategy receiving the second greatest amount of weight is one of the Lewisian signalling systems, and is

²⁷The probabilities listed do not sum to one simply because the remaining residual probability is distributed over the other twelve strategies.

the same Lewis signalling system as used by players 6 through 11. (Incidentally, it is also the strategy assigned the greatest amount of probability by player 5.) The two remaining strategies are ineffective ones for the Sender-Receiver game because they fail to distinguish between signals when in the role of Receiver. Similarly unusual outcomes (although with different probabilities) exist for the Bayesian dynamics as well. Neither dynamics provides a good model for how a population of agents might arrive at a socially optimal signalling system.

How does imitative learning fare at guiding the population to adopt one of the Lewis signalling systems? Figure 9 illustrates the outcome of one simulation using imitative learning on the lattice, where people interact with, and learn from, their eight nearest neighbours. Within a very short period of time — six iterations, in fact — the vast majority of the population has adopted one of the two Lewis signalling systems. This always happens on the lattice provided that the initial distribution of strategies contains sufficiently many of one of the two signalling systems.

“Sufficiently many” need not be that large of a number. What matters most is that a cluster of people following the same signalling system coincide. Once a block of agents has settled upon a signalling system, imitative learning will cause other agents to adopt that signalling system, causing it to spread, until the expanding region either spreads to the entire lattice or it encounters a competing region following the other signalling system. If no agents use one of the two signalling systems in the original state, the population can still manage to coordinate on a signalling system if innovation introduces new strategies into the population. The emergence of signalling systems on the lattice under imitative learning thus has dynamics similar to that of the emergence of fairness in the game of divide the dollar (Alexander and Skyrms, 1999)

On irregularly shaped networks, imitative learning can help move the population towards adopting a signalling system, but it really needs the help of spontaneous innovation as well to succeed. The reason why is that irregular social networks allow some agents to have more neighbours than others, and agents with a lot of neighbours can earn high payoffs even if they use strategies which aren't signalling systems. The high payoffs received by an agent with a lot of neighbours means that imitative learning will lead his neighbours to adopt his strategy, even though it really isn't very good for the Sender-Receiver game.

When spontaneous innovation occurs, these groups of poor communicators who are supported by a single well-connected individual can be replaced by mutants who employ a signalling system. The process is more heavily dependent upon the appearance of innovative strategies at the right place, at the right time, than on the lattice²⁸ but, even so, random networks can converge to signalling

²⁸The regularity of lattices means that signalling systems can spread rapidly once they get

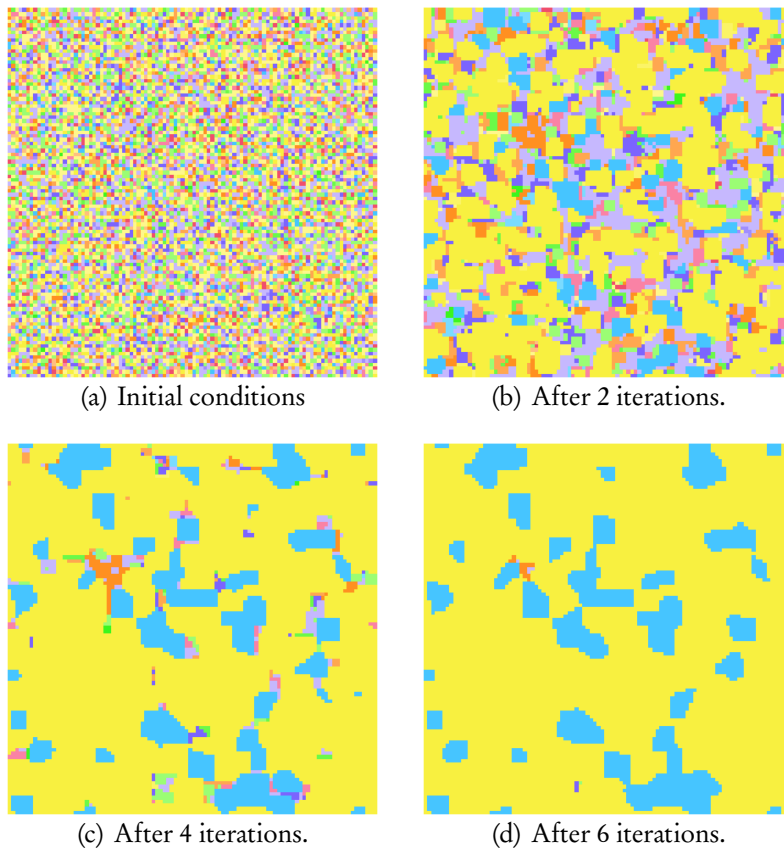


Figure 9: Imitation in the Sender-Receiver game leads to the evolution of regions where players coordinate upon one of the two Lewisian signalling systems.




Figure 10: Imitation on a random graph, with mutation.

systems in rather short order. Figure 10 illustrates the first six iterations of imitative learning on a random graph containing 25 agents. (The mutation rate was chosen so as to introduce approximately two innovative strategies into the population each iteration.) As can be seen, after six iterations the population does not seem to be converging to a signalling system. Nevertheless, the population of that simulation converged to one of the Lewis signalling systems within 50 iterations.

4. Conclusion.

There are many methods individuals might use to engage in social deliberation. In this paper, we have considered three: a variant of the Nash-Brown-von Neumann dynamics, a variant of Bayesian updating, and a form of imitative learning known as “Imitate the Best”. When the results of these three methods are compared in the Centipede game and Lewis Sender-Receiver games, the one which agrees best with what actual people do is imitative learning. In the Centipede game, only imitative learning converges to a distribution of actions in the Centipede game that tends toward the socially optimal outcome. (Both the Nash and Bayesian dynamics tend towards the traditional game-theoretic solution.) In Sender-Receiver games, only imitative learning, with innovation, is generally compatible with the population arriving at socially optimal signalling systems.

established in a local cluster.

(Both the Nash and Bayesian dynamics can cause some individuals to get stuck using inefficient strategies.)

Although imitation is not the sole method used by people to engage in social deliberation, and it may not work for all problems, it does provide a reasonably effective way of balancing the competing aims of individual maximization with social optimality. Aside from what we have seen here with the Centipede game and a simple Sender-Receiver game, imitative learning can also support cooperative behavior in the Prisoner's Dilemma, trusting behavior in the Stag Hunt, and fair division in divide-the-dollar (Alexander, 2007). If Gabriele Tarde was incorrect in asserting that society is imitation, it nevertheless is true that imitation supports many behaviours which are central to our social existence.

A. Definition of the Nash and Bayesian dynamics

Suppose that, as in section 1, we have a two-person noncooperative game. Let $\vec{p}_{\text{Row}}(t) = \langle p_1, \dots, p_n \rangle$ denote the state of uncertainty of Row, and $\vec{q}_{\text{Col}}(t) = \langle q_1, \dots, q_n \rangle$ denote the state of uncertainty for Column. Let $\text{EU}_{\text{Row}}(i, t)$ denote the expected utility of action i for Row at time t . The expected utility of the status quo for Row at time t is defined as follows:

$$\text{ESQ}_{\text{Row}}(t) = \sum_{i=1}^n p_i \text{EU}_{\text{Row}}(i, t).$$

Given these definitions, we may define the *covetability* of action j for Row at time t is:

$$\text{Cov}_{\text{Row}}(j, t) = \max \left(0, \text{EU}_{\text{Row}}(j, t) - \text{ESQ}_{\text{Row}}(t) \right).$$

Similar definitions can be made for Column. In the following, I shall omit the explicit reference to either Row or Column.

The Nash dynamics states that an individual will modify his state of uncertainty according to the rule

$$p_i(t+1) = \frac{k \cdot p_i(t) + \text{Cov}(i, t)}{k + \sum_{j=1}^n \text{Cov}(j, t)}$$

where $k > 0$ is an "index of caution" which measure how quickly individuals will adjust their probability distributions in a single revision.

The Bayesian dynamics takes the slightly different form:

$$p_i(t+1) = p_i(t) + \frac{1}{k} \cdot p_i(t) \cdot \frac{\text{EU}(i, t) - \text{ESQ}(t)}{\text{ESQ}(t)}$$

where, again, $k > 0$ provides an index of caution reflecting how rapidly the distribution changes in a single revision. Skyrms (1990, pg. 36–38) explains the connection between the above formula and Bayes’ theorem. Before applying the Bayesian dynamics, one must first transform the payoff matrix so that the lowest payoff is 0 and the greatest payoff is 1; without the transformation, the result will not necessarily be a probability distribution.

References

- Alexander, J. McKenzie (2007). *The Structural Evolution of Morality*. Cambridge University Press. 4, 22
- (2009). “Local Interaction and the Dynamics of Rational Deliberation”. *Philosophical Studies* Forthcoming. 4
- Alexander, Jason and Brian Skyrms (1999). “Bargaining with Neighbors: Is Justice Contagious?” *Journal of Philosophy* 96, 11: 588–598. 19
- Fey, Mark, Richard D. McKelvey, and Thomas R. Palfrey (1996). “An experimental study of constant-sum centipede games”. *International Journal of Game Theory* 25: 269–287. 7
- Gigerenzer, Gerd and Reinhard Selten, eds. (2001). *Bounded Rationality: The Adaptive Toolbox*. MIT Press. 6
- Hollis, Martin (1994). *The Philosophy of Social Science: An Introduction*. Cambridge University Press. 8
- Latour, Bruno (2002). “Gabriel Tarde and the end of the social”. In Patrick Joyce, ed., *The Social in Question*, Routledge, pp. 117–132. 2
- Lehrer, K. and C. Wagner (1981). *Rational Consensus in Science and Society*. Dordrecht:Reidel. 5
- Lewis, David K. (1969). *Convention: A Philosophical Study*. Harvard University Press.
- McKelvey, Richard D. and Thomas R. Palfrey (1992). “An experimental study of the centipede game”. *Econometrica* 60, 4: 803–836. 7, 14
- Millet, Jean (1970). *Gabriel Tarde et la philosophie de l’Histoire*. Paris: Vrin. 2
- Rosenthal, R. (1981). “Games of Perfect Information, Predatory Pricing, and the Chain Store”. *Journal of Economic Theory* 25: 92–100. 6

- Skyrms, Brian (1990). *The Dynamics of Rational Deliberation*. Harvard University Press. 4, 23
- (2008). “Evolution of signaling systems with multiple senders and receivers”. *Phil. Trans. R. Soc. Lond. B* Working paper. 15
- Smead, Rory (2008). “The Evolution of Cooperation in the Centipede Game with Finite Populations”. *Philosophy of Science* 75: 157–177. 15
- Tarde, Gabriel (1903). *The Laws of Imitation*. Henry Holt and Company. Translated from the second French edition by Elsie Clews Parsons. Original French edition published in 1890. 2
- Vanderschraaf, Peter and J. McKenzie Alexander (2005). “Follow the Leader: Local Interactions with Influence Neighborhoods”. *Philosophy of Science* 72: 86–113. 17