Evolution, morality and the theory of rational choice

J. McKenzie Alexander

Within moral philosophy, much work has been done to show that the constraints imposed by morality are either compatible with, or derivable from, the demands of rationality.[1] Whether a serious problem exists in reconciling the two largely depends on one's exact conception of both morality and rationality. Within the Kantian tradition, for example, no such problem exists as the moral law derives from considerations of practical reason. Within the tradition of contemporary economic thought (and, to some extent, contemporary social scientific thought), the problem looms large. Rationality conceived of in individualistic, or instrumental, terms may conflict with the demands of morality whenever what lies in the perceived self interest of the individual conflicts with what morality requires.

Such conflict, of course, may not exist for all individuals, even if we equate rationality with instrumental rationality. A moral agent's individual preferences are aligned with the demands of morality such that when he acts to maximise expected utility, he will also act morally.[2] Only imperfect moral agents will feel the conflict with the demands of morality, for at times what lies in their self interest will be at odds with what morality requires.

The reason that the behaviour of moral agents, when they act to maximise expected utility, agrees with the requirements of morality is because the preferences of moral agents have been explicitly shaped so as to conform with the demands of morality. *Given* a set of moral rules, education and habituation can cause a individual's preferences to become compatible with that set of rules. This is nothing new: Aristotle notes this relationship between education, habituation, and morality in the *Nichomachean Ethics*.

Yet this raises the following question regarding the *emergence* of morality: in a society or population existing in a pre moral state, where individual preferences may range freely over outcomes and actions with no restrictions whatsoever, why would individual's preferences tend towards the shape they take when moulded by morality, before morality exists? It is with this question that the conflict between morality and instrumental rationality becomes especially acute. Morality, at times, requires that individuals refrain from actions which would materially benefit them in order to bring about a fair division of resources, or to bring about a common good (Think of Garrett

[1] By »morality« I am referring to our pre theoretic sense of what we ought to do.

[2] It is a mistake to think that the framework of expected utility theory requires agents to be selfish or that it precludes the possibility of other-directed preferences. Agents are »self interested« in the sense that what motivates them to act are simply the interests that they themselves possess. Those interests may directly concern the well-being of others and may be altruistic.

Generelle Anmerkung:

Bitte schauen Sie sich den gesammten Text nochmal genau an um eventuelle Abweichungen vom Manuskript auszuschließen.

Hardin's famous example of the tragedy of the commons.) If people in a pre moral state only care about their own material gain, they will not want to act so as to bring about the common good, or to bring about a fair division. But presumably people must first *want* to act to bring about a common good, or a fair division, before they adopt or formulate a morality which tells them that they *ought* to act to bring about a common good, or a fair division.[3] Why would people, in a pre-moral state, possible adjust their individual preferences so as to make them compatible with the requirements of morality, if all they care about is maximising their own personal gain?

In what follows, I will take some steps towards showing that the apparent conflict between the demands of instrumental rationality and morality can be mitigated by conceiving of our rational capacity and moral outlook as evolutionary products. As both instrumental rationality and morality are normative theories specifying what ideal agents *ought* to do, conflict between the two normative theories should not surprise us. Unless the two normative theories are carefully developed so as to be compatible by design, conflict should be expected.[4] I suggest that behaviours which evolution favours in »rational choice« problems and the behaviours which evolution favours in »moral choice« problems will necessarily be compatible, as both serve to maximise fitness. What I will show is that, in the Ellsberg decision problem, the Centipede game, and the Ultimatum game, the behaviours selected by evolution accord quite well with what our intuitions tell us to do. This is of particular interest because, in all three decision problems, our intuitions flagrantly violate what the traditional theory of rational choice recommends.

1. Evolution and the Ellsberg Paradox

Consider the following decision problem, due to Ellsberg (1961): an urn contains a mix of ninety red, black, and yellow balls. Thirty of the balls are red and the remaining sixty are black or yellow, although the exact composition is not known. (It could be the case that all sixty are black or that all sixty are yellow, or any distribution in between.) You are then presented with two choices of gambles:

[3] This assumes a very thin conception of morality. Generally speaking, I shall be thinking of morality as nothing more than a set of particular social norms of a group, where these social norms are a product of historical accident combined with what the society believes is in their collective interest. »J. Alexander, *The Structural Evolution of Morality*, 238–266« elaborates upon this conception of morality.

[4] And notice that theories which strive, by design, to eliminate conflict between what rationality requires and what morality demands, such as Kantian philosophy and Utilitarianism, typically suffer on the grounds that they get the moral requirements wrong, the rational requirements wrong, or both. (By »wrong« what I mean is that the requirements conflict, sometime strongly, with our pre-theoretic intuitions regarding what we should do.) The Kantian conception of rationality looks very strange from a contemporary point of view in uenced (some may prefer »contaminated«) by economics and modern decision theory. Likewise, many of the Kantian moral requirements strike us as unreasonable. (You should never lie under any circumstance? You should do what morality requires even if the world will end as a result?) The same can be said for Utilitarianism. (Consider Parfit's Repugnant Conclusion and the attempts to deal with it.)

1.	Gamble A Win £100 if you draw red.	or	Gamble B Win £100 if you draw black.
2.	Gamble C Win £100 if you draw red or yellow.	or	Gamble D Win £100 if you draw black or yellow.

Given these options, what would you select in case 1 and case 2?

If you are like most people, you would pick Gamble A in the first case and Gamble D in the second case. This choice pattern violates the Sure-Thing principle from Savage's theory of expected utility, and one can easily show that there is no coherent utility function and assignment of probabilities attributable to individuals which licences this choice.[5] Given this, how can we explain why individuals are disposed to choose this way?

To begin, note that *ambiguity aversion* most likely in uences people's behaviour in the Ellsberg paradox.[6] Decisions made as a result of ambiguity aversion can still be explained in terms of expected utility provided that one adopts a nonstandard theory of expected utility. For example, Schmeidler[7] suggests replacing the independence axiom with one regarding co-monotonicity of preferences.[8] From this, he develops a theory which has unique nonadditive probabilities and a von Neumann/Morgenstern utility

5 Suppose that there was such a function $u(\cdot)$. Let $\Pr(R)$, $\Pr(B)$ and $\Pr(Y)$ denote the probability assigned by the person to drawing a red, black, and yellow ball, respectively. Choosing Gamble A over Gamble B implies

$$\Pr(R) \cdot u(100) + (1 - \Pr(R)) \cdot u(0) > \Pr(B) \cdot u(100) + (1 - \Pr(B)) \cdot u(0)$$

and choosing Gamble D over Gamble C implies

$$(\Pr(R) + \Pr(Y)) \cdot u(100) + (1 - \Pr(R) - \Pr(Y)) \cdot u(0) < (\Pr(B) + \Pr(Y)) \cdot u(100) + (1 - \Pr(B) - \Pr(Y)) \cdot u(0).$$

Inspection of the second inequality shows that all terms containing »$\Pr(Y)$« can be eliminated from it, with the resulting inequality contradicting the first.

6 Let B denote the number of black balls in the urn and Y the number of yellow balls in the urn. Then the chance of an individual winning the respective gambles (which he does not know) are as follows:

1.	Gamble A $\frac{30}{90}$	Gamble B $\frac{B}{90}$
2.	Gamble C $\frac{30+Y}{90}$	Gamble D $\frac{60}{90}$

If individuals wish to avoid ambiguity and *this alone* determines their choice, then picking Gamble A in case 1 and Gamble D in case 2 is the only option.

7 »D. Schmeidler, ›Subjective Probability and Expected Utility without Additivity‹, 571–587.«

8 Let s and t be states of the world, and let f and g be two acts. Then f and g are co-monotonic if it never is the case that $f(s) \succ f(t)$ and $g(t) \succ g(s)$. Schmeidler's axiom states that if f, g, and h are comonotonic acts and $f \succ g$, then $\alpha f + (1 - \alpha)h \succ \alpha g + (1 - \alpha)h$, where $\alpha \in [0,1]$.

function.[9] Alternatively, Gilboa and Schmeidler[10] take an approach in which individuals are permitted to have multiple priors. Both cases allow representation of ambiguity aversion, albeit in different ways.

Yet this raises the question: whence does ambiguity aversion originate? Let us approach this question from an evolutionary perspective, if only for the reason that no obvious *a priori* reason exists for why natural selection should be concerned with selecting for conformity to axioms of rational choice like the Sure-Thing principle. Natural selection aims to maximise fitness within a population. If environmental conditions are such that maximal fitness occurs when behaviour violates principles of rational choice, so much the worse for compliance with those principles.[11]

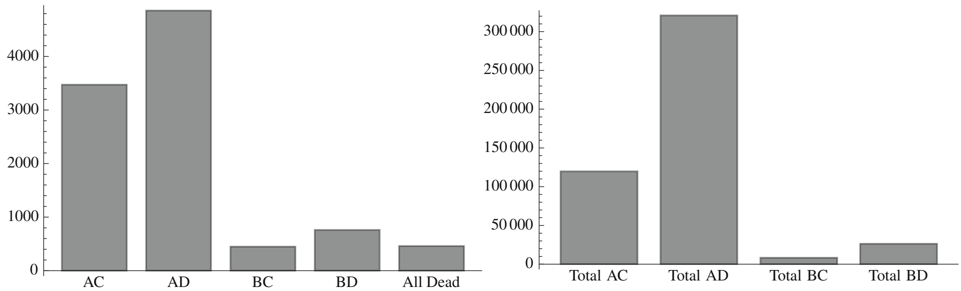

(a) Number of simulations which converged to a pure behavioural state.

(b) Overall reproductive success of the various player types.

Figure 1: Outcome of 10,000 simulations of the Ellsberg decision problem.

Consider, then, the following evolutionary model. Suppose that people are of one of four »types«: AC, AD, BC, or BD. The type of a person indicates their choice behavior in both cases of Ellsberg's decision problem. Suppose also that, at the start of each generation, Nature fills the urn with a random mix of black and yellow balls and that Nature also decides whether to present people with choice (1) or choice (2). For simplicity, let us assume that Nature fills the urn with a random mix (b,y) chosen from {(0,60), (1,59), … , (59,1), (60,0)} with each outcome equally likely. Every person in the population then draws a ball from the urn (with replacement). A »winner« has two offspring of the same type as the parent, a »loser« has no offspring.

9 It remains a nonstandard theory of expected utility because one must compute expected utilities using the Choquet integral.

10 »I. Gilboa and D. Schmeidler, ›Maxmin Expected Utility with Nonunique Prior‹, 141–153.«

11 A similar thing could be said regarding truthfulness or accuracy of representations. Natural selection would only support truthful beliefs and accurate representations of the world when those are aligned with maximising fitness. As some evolutionary explanations of religious belief suggest, these beliefs persist because their are either (a) fitnessenhancing, or (b) mis-firings of other psychological processes or mechanisms which are fitness-enhancing (see »R. Dawkins, *The God Delusion*« and »D. Dennett, *Breaking the Spell: Religion as a Natural Phenomenon*.«

schließende Klammer fehlt, nach "... Phenomenon.)"setzen?

Figure 1 illustrates the results of 10,000 simulations of this evolutionary model. The population was started in a state containing 200 individuals assigned types at random.[12] The simulation was run for 1,000 generations or until either (a) every individual died out, or (b) convergence to a population state containing only a single type occurred. (As one can see from the overall summary in figure 1(a), very few simulations failed to converge.) Over the course of the simulation, the overall number of offspring of each type was also tracked.

Strikingly, the »irrational« type AD proved most fit on two different measures. First, more populations converged to a state containing all-AD than any other competing type. Second, from the point of view of maximising the overall number of offspring, AD was more successful than any other competing type, by far (see figure 1(b)). Thus we have identified one possible reason for people's seemingly »irrational« behaviour in the Ellsberg problem: avoiding ambiguity is evolutionarily advantageous. Why? Because evolution is sensitive to the *variance* in fitness, as well as *expected* fitness.[13]

2. Evolution and the Centipede game

The Centipede game[14], first introduced by Rosenthal[15], provides a wellknown example of a decision problem where the traditional game theoretic analysis conflicts with what our »moral« intuitions suggest as the way to play.[16] Figure 2 illustrates a six-stage Centipede game. 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 occurs when both players always choose Pass. However, if one solves the game using

[12] The random types were drawn from a randomly selected distribution. The random distribution was determined via a »stick-breaking« algorithm: three random numbers were generated in the interval [0,1], determining 4 subintervals. The length of the *i*th subinterval was the probability assigned to the *i*th type of player. This produces an unbiased sampling from the space of possible distributions.

[13] By construction, the chance of willing Gamble B when the urn is filled in the manner specified is exactly $\frac{1}{3}$, the same as winning Gamble A. (This follows trivially since Pr(Black) = $\frac{1}{61}\sum_{i=0}^{60}\frac{i}{60}=\frac{1}{3}$.) Similarly, the chance of winning Gamble D is the same as winning Gamble C. From this we know that the fitness advantage conferred to AD over the other strategies has nothing to do with the expected number of offspring (which is the same for all four types) but rather the variance in the number of offspring (which differs amongst all four types).

[14] Portions of this section were drawn from »J. Alexander, ›Social Deliberation: Nash, Bayes, and the Partial Vindication of Gabriele Tarde‹.«

[15] »R. Rosenthal, ›Games of Perfect Information, Predatory Pricing, and the Chain Store‹, 92–100.«

[16] I include the term ›moral‹ in scare quotes because what really underlies the conflict is the competing pull between achieving the socially optimal outcomes and achieving the individually rational outcome. Yet the socially optimal outcome is, in many cases, what morality requires, as noted in the introduction, and so the centipede game may, in some instances, represent problems where instrumental rationality conflicts with morality.

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.[17]

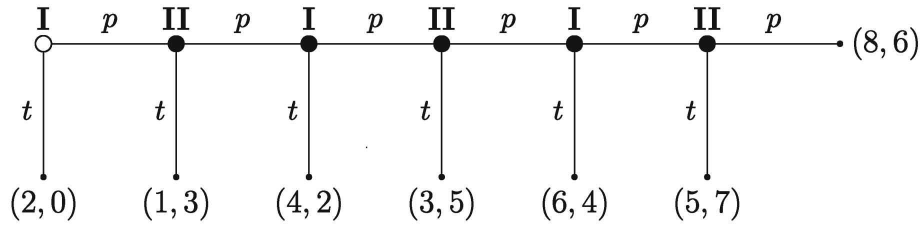

Figure 2: A six-stage Centipede game.

Backwards induction recommends player I take on the first move. Yet many people's intuitions suggest that what the players should do is choose *Pass* for a number of moves, potentially even to the very end of the game. These intuitions are borne out by experiment, as McKelvey and Palfrey[18] report. In a six-stage centipede game, they found that only 1% of the players choose *Take* on the first move. Moreover, if the game should happen to reach the final stage, 15% of the time the last player would chooses *Pass*, thereby playing a dominated strategy but, at the same time, producing the socially optimal outcome.

As before, let us consider this from an evolutionary perspective. One important difference between the Centipede game and the Ellsberg problem is that the Centipede game involves *strategic* choice whereas the Ellsberg problem only involved *parametric* choice.[19] As such, we need to adopt a framework capable of modeling the evolutionary dynamics of strategic problems.

Alexander[20] has advocated the use of *local interaction models* as a way of representing problems of strategic choice in a socially structured context. Let us adopt that framework here for the Centipede game. More formally, suppose that we have a population P = {1,…, n} of players and let $G = \langle P, E \rangle$ be a graph representing a binary relation of some social importance, such as »being an acquaintance of«, »being related to«, and so on. If two players lie on an edge $e \in E$, they are said to be

17 Consider 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.

18 »R. McKelvey and T. Palfrey, ›An experimental study of the centipede game‹, 803– 836.«

19 In the Ellsberg problem, individuals play a »game against Nature«, in that Nature fills the urn and chooses which case to present people with. Nature, though, does not attempt to second-guess the individuals as she does not care whether anyone wins or loses. In problems of strategic choice, the interests of the players may be completely misaligned (as in zero-sum games), partially aligned (as in the case here) or completely aligned (as in pure coordination problems).

20 »J. Alexander, *The Structural Evolution of Morality*, 25–52.«

neighbours. The set of neighbours for a given player determines whom he interacts with when he plays a game, and it also determines whom he learns from when he undergoes a process of strategic revision. Figure 3 illustrates one such graph for a population of twenty-three players.

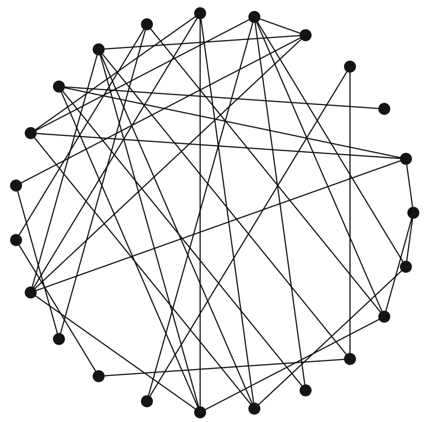

Figure 3: A graph representing a binary relation on a population of 23 individuals.

The evolutionary model proposed divides each generation, or iteration, into two stages. In the first stage, a two-player game occurs for each edge in the graph. For each edge, a fair coin is tossed to determine who plays as Player I and who plays as Player II. (This means that a single player may, in a given generation, play the Centipede game both as Player I and as Player II.) Each player in the population receives a total score equaling the sum of his payoffs from each individual game.

In the second stage, players undergo a process of strategic revision to determine whether they need to change their strategy. Rather than model individuals as perfectly rational agents capable of extremely complex maximizing calculations, I assume that individuals are boundedly rational, in the sense of Gigerenzer[21] and Gigerenzer and Selten[22]. That is, I assume that individuals determine what type of new strategy to adopt using a simple rule-governed heuristic. There are a number of heuristics worthy of consideration, but one which stands out as particularly interesting is one known as *Imitate-the-Best*. According to this heuristic, an individual *i* looks at all of his neighbours and compares the score he received at the end of the first stage with the score received by his neighbours at the end of the first stage. Let *s* denote the maximal score earned by all of his neighbours, and let *S* denote the set of strategies used by his maximally-scoring neighbours. If *s* is greater than *i*'s score, then *i* will

21 »G. Gigerenzer, P. Todd and the ABC Research Group, *Simple Heuristics That Make Us Smart*.«

22 »G. Gigerenzer and R. Selten, *Bounded Rationality: The Adaptive Toolbox*.«

select a strategy chosen at random from *S* and adopt it. If *i*'s score is greater than or equal to *s*, he will use the same strategy as before in the next round of play. After every individual has undergone a process of strategic revision, the generation ends and the next generation begins.

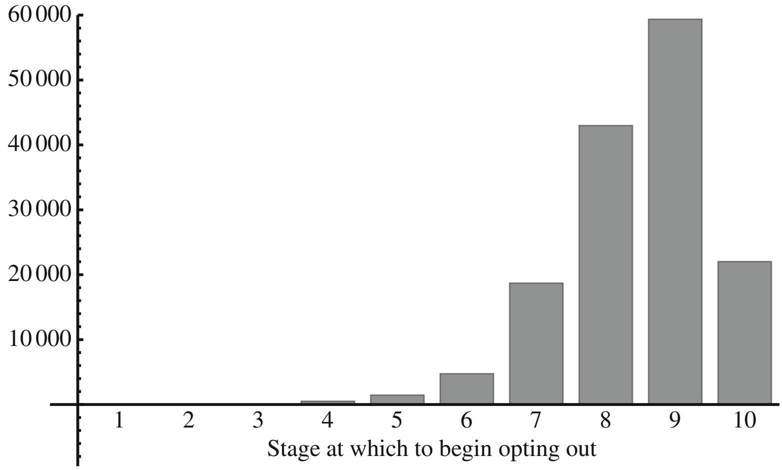

Figure 4: Simulation results for 1,000 simulations run using a 10-stage Centipede game. The *y*-axis counts the number of strategies present once the simulation has arrived at a fixed point (every individual following the same strategy) or run for 200 iterations.

What happens when individuals play the Centipede game in this local interaction setting? We can investigate this question by simulation. In order to avoid biasing the simulation results by fixing the topology of the graph, or by fixing the initial allocation of strategies to the population, we can run a number of simulations initialized randomly and tabulate the outcome.[23]

Figure 4 illustrates the outcome of one series of 1,000 simulations. In these simulations, the population consisted of 150 agents and a connected graph with a 3% edge probability.[24] (The requirement of connectivity was to ensure that the population was not broken into separate subgroups incapable of in uencing one another.) Strategies were initially assigned to players at random from a randomly chosen distribution.

Figure 4 nicely illustrates, as before, that evolution may not select the solution selected by backwards induction. Instead of each player electing to opt-out at the earliest

[23] As described above, there is no mechanism for the *introduction* of new strategies into the population. Since every agent must begin by following *some* strategy, this means that the set of possible strategies to which the simulation might converge (if it does) is determined by the initial conditions. Thus it is important to consider not only randomly generated graphs (to avoid results excessively dependent upon the graph topology) but to consider randomly chosen initial conditions (to avoid results excessively dependent upon the original assignment of strategies).

[24] For a graph containing 150 agents, there are $\binom{150}{2} = \frac{150 \cdot 149}{2} = 11{,}175$ possible edges. With a 3% edge probability, that means each graph contained approximately 335 edges.

possible stage, we find that players tend to opt-out at later stages. The skew of the distribution rather far towards the right seems to accord with our intuitions about the »right« way to play the Centipede game far better than with the tradition game-theoretic analysis.

Moreover, recall that McKelvey and Palfrey[25] found that individuals arrived at the socially optimal outcome approximately 15% of the time. In this model, the socially optimal outcome occurred when individuals chose *Take* at stage 10. (Choosing *Pass* at the last stage was not an option.) Here, approximately 20,000 individuals followed the socially optimal outcome at the end of the simulation. Since there were 1,000 simulations in total, with 150 individuals in each simulation, this means that the socially optimal outcome was obtained approximately $\frac{20{,}000}{150 \cdot 1{,}000} \approx 13{,}3\%$ of the time.

3. Evolution and the Ultimatum game

The Centipede game is not the only game where considerable disagreement has been found between the outcomes of experiment and the game-theoretic solution. Another famous example is known as the »Ultimatum game«, which elicits intuitions concerning fair outcomes in asymmetric situations. In the Ultimatum game, two players are randomly assigned the roles of »Proposer« and »Receiver«. The Proposer is given a fixed sum of money C and must offer some amount to the Receiver. The offer may consist of any feasible amount in the interval $[0, C]$. The Receiver, upon notification of the offer, may choose either to accept the offer or reject it. If the offer is accepted, then the Receiver gets the amount offered and the Proposer keeps the remaining amount; if the offer is rejected, though, neither player receives anything.

One can easily see that if players are instrumentally rational and the game is only played once, then the game-theoretic solution is for the Proposer to offer the smallest positive feasible amount, which the Receiver will accept. Assuming that something is better than nothing, the Receiver will never reject an offer c provided that $c > 0$. The Proposer, knowing this, will offer the smallest amount c possible. When played with real money, say £10, with offers restricted to £1 increments, this means that the Proposer will offer £1, which the Receiver will accept, leaving the Proposer with £9 for himself.

In a seminal paper, Güth, Schmittberger and Schwarze[26] found that when people play the Ultimatum game, their behaviour lies quite far removed from the game-theoretic solution. In fact, the *modal* offer for Proposers was to offer around 40% of the money (where this offer would generally be accepted). Moreover, when Proposers tried to take advantage of their asymmetric position by offering only 20% of the money, this offer would generally be rejected, even though the amount offered was not negligible.

25 »R. McKelvey and T. Palfrey, ›An experimental study of the centipede game‹, 803– 836.«

26 »W. Güth, R. Schmittberger and B. Schwarze, ›An Experimental Analysis of Ultimatum Bargaining‹, 367–388.«

Since the time of Güth et al.'s original experiment, a very large number of experiments on a number of variations of the Ultimatum game have been performed.[27]

The results of Henrich et al.[28] suggest that one cannot escape appealing to social norms or other cultural forces in explaining human behaviour in the Ultimatum game. What I want to illustrate in this section is simply how local interaction models of the Ultimatum game can generate a range of behaviours with are, to some extent, compatible with the experimental findings. If human behaviour in the Ultimatum game is constrained by people's beliefs about what is *fair* in asymmetric resource allocation problems, this then shows that the outcome of a cultural evolutionary process is, to some extent, compatible with observed human behaviour. In any case, we find a much better agreement of the *variation* of human behaviour if we consider it as the outcome of an evolutionary process then if we consider it as generated by the decisions of a perfectly rational agent.

Consider the following model, which is similar to that of Nowak, Page and Sigmund[29] but with different evolutionary dynamics. Suppose that each player has a strategy consisting of two components $p, q \in [0, 1]$, where p denotes the proportion of C offered when in the role of a Proposer and q denotes the acceptance threshold. When a player acts as a Proposer, she always offers p; when a player acts as a Receiver, she will accept any offer r provided that $q \leq r$.

27 »R. Thaler, ›Anomalies: The Ultimatum Game‹, 195–206« offers an early survey of the experimental results, albeit now rather dated. See »C. Camerer and R. Thaler, »Anomalies: Ultimatums, Dictators and Manners«, 209–219« for a follow-up and »C. Bicchieri, *The Grammar of Society*, 100–139« for a more recent discussion of experimental work on the Ultimatum game. »J. Henrich, R. Boyd, S. Bowles, C. Camerer, E. Fehr and H. Gintis, *Foundations of Human Sociality*« cover a number of ultimatum-game type experiments in »small-scale« societies, noting how considerable deviations exist as compared with previously reported results concentrating on Western societies. For example, the Au and Gnau of New Gineau reject both unfair and hyperfair offers with approximately equal frequency. (A hyperfair offer is one where the Proposer offers more than half of C to the Receiver.)

28 »J. Henrich, R. Boyd, S. Bowles, C. Camerer, E. Fehr and H. Gintis, *Foundations of Human Sociality*.«

29 »M. Nowak, K. Page and K. Sigmund, »Fairness versus Reason in the Ultimatum Game«, 1773–1775.«

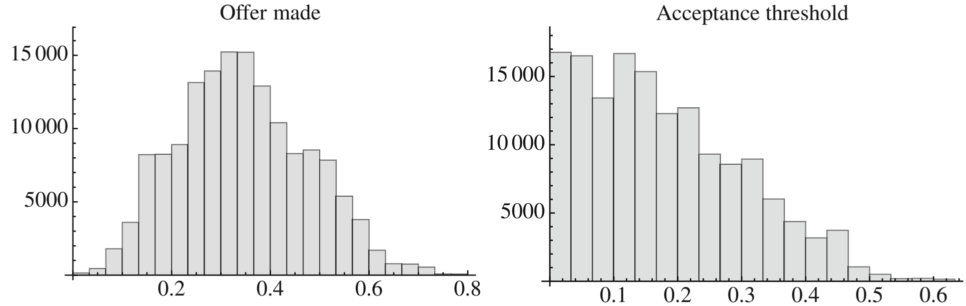

(a) Histogram of offers made when in the role of Proposer. (The x-axis indicates the fraction of C offered.)

(b) Histogram of player's acceptance threshold when in the role of Receiver. (The x-axis indicates the amount of cake which the player must be offered in order to accept.)

Figure 5: Results from 1,000 simulations of the Ultimatum game on random connected networks of 150 agents with an edge probability of 3%.

As in the case of our Centipede game simulations, we consider populations of 150 players situated on a connected social network with an edge probability of 3%. During the first stage of each generation (the interaction stage), for each edge in the network a coin toss determines which of the two players on that edge is assigned the role of Proposer and Receiver. As before, this means that the same player may, in a given generation, act as both the Proposer and Receiver. During the second stage of each generation (the revision stage), players use *Imitate-the-Bes*t to adopt the strategy of the maximally-scoring player in their neighbourhood. When a player imitates a strategy of one of his neighbours, he copies both the offer type p and the acceptance threshold q, though he only saw one of these components used in his last interaction.

Figure 5 illustrates the outcome of 1,000 simulations on random social networks from randomly chosen initial conditions.[30]One curious aspect worth noting is that, although the acceptance thresholds of the players appears to conform somewhat to what one would expect from rational agents (in that low acceptance thresholds appeared more frequently than high ones), the distribution of offer types does not. The frequency of offer types is highest around the offers of 0.3 to 0.375 of C. Perhaps the two most noteworthy aspects about the simulation outcomes are the following:

1. Considerable variation exists among the convergent states of the simulation. No one offer type overwhelming dominates among the results. This qualitative result agrees with what we know from experimental data.

[30] The initial strategy (p, q) of each player was drawn at random from the unit square under a uniform distribution.

2. The vast majority of offer types lie below the 50–50 split, with some cases converging to hyperfair offers. This qualitative result agrees with the findings of Henrich et al.[31]

Once again, we find that evolution yields outcomes more in agreement with actual human behaviour than the predictions of traditional game theory.

4. Conclusion

Human behaviour in experimental settings frequently violates the predictions of decision theory and game theory in interesting ways. In this paper I have attempted to show that, for the Ellsberg decision problem, the Centipede game, and the Ultimatum game, considering such problems from an evolutionary perspective yields models with quite good qualitative agreement with the experimental outcomes. Obviously these models only provide a first step towards modeling decision making in socially structured contexts, but I believe two points are worth noting. First, we do not need complicated models involving a plethora of social factors to achieve qualitative agreement with observed human behaviour; instead, relatively simple models suffice. Second, that the outcomes of an evolutionary process tend to agree both with what strikes us as intuitively »rational« or »fair«. Thus, as I stated in the introduction, this suggests that one way of reconciling the apparent competing demands of rationality and morality may be to consider *both* from an evolutionary perspective.

References

Alexander, J. McKenzie: *The Structural Evolution of Morality.* Cambridge University Press, London 2007.

— : »Social Deliberation: Nash, Bayes, and the Partial Vindication of Gabriele Tarde«, in: *Episteme* Forthcoming 2009.

Bicchieri, Cristina: *The Grammar of Society.* The Nature and Dynamics of Social Norms, Cambridge University Press, London 2005.

Camerer, Colin and Richard H. Thaler: »Anomalies: Ultimatums, Dictators and Manners«, in: *Journal of Economic Perspectives* 9, 2 (1995) 209–219.

Dawkins, Richard: *The God Delusion.* Houghton Mif in Harcourt 2006.

Dennett, Daniel C.: *Breaking the Spell: Religion as a Natural Phenomenon.* Viking Adult 2006.

Ellsberg, Daniel: »Risk, Ambiguity, and the Savage Axioms«, in: *Quarterly Journal of Economics* 75, 4 (1961) 643–669.

[31] »J. Henrich, R. Boyd, S. Bowles, C. Camerer, E. Fehr and H. Gintis, *Foundations of Human Sociality*, 19.«

Gigerenzer, Gerd and Reinhard Selten: *Bounded Rationality*. The Adaptive Toolbox, MIT Press, 2001.

Gigerenzer, Gerd, Peter M. Todd, and the ABC Research Group: *Simple Heuristics That Make Us Smart*. Oxford University Press, 1999.

Gilboa, I. and D. Schmeidler: »Maxmin Expected Utility with Nonunique Prior«, in: *Journal of Mathematical Economics* 18 (1989) 141–153.

G uth, Werner, Rolf Schmittberger, and Bernd Schwarze: »An Experimental Analysis of Ultimatum Bargaining«, in: *Journal of Economic Behavior and Organization* 3 (1982) 367–388.

Hardin, Garrett: »The Tragedy of the Commons«, in: *Science* 162 (1968) 1243–1248.

Henrich, Joseph, Robert Boyd, Samuel Bowles, Colin Camerer, Ernst Fehr, and Herbert Gintis: *Foundations of Human Sociality*. Economic Experiments and Ethnographic Evidence from Fifteen Small-Scale Societies, Oxford University Press 2004.

McKelvey, Richard D. and Thomas R. Palfrey: »An experimental study of the centipede game«, in: *Econometrica* 60, 4 (1992) 803–836.

Nowak, Martin A., Karen M. Page, and Karl Sigmund: »Fairness Versus Reason in the Ultimatum Game«, in: *Science* 289, 5485 (2000) 1773–1775.

Parfit, Derek: *Reasons and Persons*. Oxford University Press 1984.

Rosenthal, R.: »Games of Perfect Information, Predatory Pricing, and the Chain Store«, in: *Journal of Economic Theory* 25 (1981) 92–100.

Schmeidler, David: »Subjective Probability and Expected Utility without Additivity«, in: *Econometrica* 57, 3 (1989) 571–587.

Thaler, Richard H.: »Anomalies: The Ultimatum Game«, in: *Journal of Economic Perspectives* 2, 4 (1988) 195–206.