SCIENTIFIC REP🌼RTS

The Evolution of Generosity in the Ultimatum Game

Arend Hintze[1,2,3] & Ralph Hertwig[4]

When humans fail to make optimal decisions in strategic games and economic gambles, researchers typically try to explain why that behaviour is biased. To this end, they search for mechanisms that cause human behaviour to deviate from what seems to be the rational optimum. But perhaps human behaviour is not biased; perhaps research assumptions about the optimality of strategies are incomplete. In the one-shot anonymous symmetric ultimatum game (UG), humans fail to play optimally as defined by the Nash equilibrium. However, the distinction between kin and non-kin—with kin detection being a key evolutionary adaption—is often neglected when deriving the "optimal" strategy. We computationally evolved strategies in the UG that were equipped with an evolvable probability to discern kin from non-kin. When an opponent was not kin, agents evolved strategies that were similar to those used by humans. We therefore conclude that the strategy humans play is not irrational. The deviation between behaviour and the Nash equilibrium may rather be attributable to key evolutionary adaptations, such as kin detection. Our findings further suggest that social preference models are likely to capture mechanisms that permit people to play optimally in an evolutionary context. Once this context is taken into account, human behaviour no longer appears irrational.

The study of the one-shot anonymous ultimatum game (UG) has helped researchers to understand what influence social preferences[1] have on human choice. In this game, two players—a *proposer* and a *responder*—are tasked with an economic decision[2]. First, the *proposer* suggests how to split a resource. This resource, provided by a third party, is set to 1.0. The offer defines the fraction that the proposer is willing to give to the responder. If the offer meets the responder's demands (is accepted), the resource will be split accordingly. If the offer is rejected, neither player will receive anything. Our focus is on the standard version of the UG: a one-shot game between anonymous parties. A player's strategy is defined as a vector (*offer*, *demand*). Note that these values are not probabilities but fractions within the range of [0, 1]. Therefore, strategies play deterministically, and the strategy space is theoretically infinite. The rational and most optimal strategy in this game differs markedly from the strategy observed in humans. This disparity has been explained by showing that humans take other elements, such as the social context, into account when making decisions. Here, we show that allowing players to discern between opponents that are kin and opponents that are not kin leads to the evolution of strategies that resemble those observed in humans—even if kinship detection is imperfect.

The UG has been extensively studied computationally in contexts such as spatial situations[3,4], in games played on networks[5,6], in closely related games[7,8], and in the absence of complete information[9]. As we were interested in the *evolution* of strategies, we used a computational evolutionary model[10–12] that defined a population of strategies as a set of points within the strategy space. We randomly selected a set of strategies to form the population of the first generation. These selected strategies all played against each other, and not against every possible strategy— much as an organism does not interact with all other possible organisms, but only with those alive at the time. The next generation was formed from those strategies that earned a higher payoff playing the UG. Specifically, the probability of an organism having offspring was defined by the payoff it received divided by the total payoff the entire population received (see Methods and Supplementary Information for a full description of the model). However, mutations could move the points defining the offspring's strategy around in this space, providing a new set of opponents. How this set of points moved through this space over time approximated how evolution would act on organisms whose existence depended solely on playing the UG.

[1]BEACON Center for the Study of Evolution in Action, Biomedical and Physical Sciences Building, 567 Wilson Road Room 1441, East Lansing, MI 48824, USA. [2]Department of Integrative Biology, Michigan State University, 288 Farm Lane RM 203, East Lansing, MI 48824, USA. [3]Department of Computer Science and Engineering, Michigan State University, 428 South Shaw Lane RM 3115, MI 48824, USA. [4]Center for Adaptive Rationality, Max Planck Institute for Human Development, Berlin, Germany. Correspondence and requests for materials should be addressed to A.H. (email: hintze@msu.edu)

Our goal is to appropriately model the evolution of strategies playing the UG and to compare them with human behaviour. Unfortunately, the strategy vectors of humans (*offer*, *demand*) cannot be derived instantly from observing them play. Instead, offer distributions are reported in terms of empirical observations, and demands are presented as acceptance rates, conditioned on the offer. Thus, we first translate observations of human behavior into a strategy vector that can be directly compared with our findings obtained from computationally evolving strategies to play the UG. We then endowed these strategies with the ability to detect kinship relations. By comparing the trajectories of strategies playing the UG with and without the ability to detect kin, we were able to study the effect of kinship recognition on the optimality of evolved strategies.

From the point of view of economic rationality, however, how a set of strategies moves through space is not as important as how individuals optimize their own payoff. Specifically, the success of a strategy depends on its ability to optimize the offer in such a way that the smallest acceptable fraction is handed over. At the same time, demand has to be optimized such that the responder receives as much as possible and does not become too renunciative. The Nash equilibrium[13,14] for the UG is to offer the lowest possible amount that is not zero, while at the same time accepting any non-zero offer (*offer* ~ 0.0, *demand* ~ 0.0). This can be directly derived when considering the subgame perfect equilibrium[15,16]: The responder has to choose between the present offer or nothing. Therefore, a responder with purely self-regarding preferences (i.e., maximization of self-interest) will accept any offer that is larger than zero, because anything is better than nothing. The proposer knows this, and can thus maximize her payoff by making the smallest possible non-zero offer. Offers optimized according this rationale should approach zero.

However, the strategy vector revealed in human play deviates[2,17–23] from the rational Nash equilibrium (see the initial study by Güth *et al.*[2]; for a meta-analysis, see Spiegel *et al.*[24]) on three dimensions. First, the human strategy is more *generous*[25]; that is, proposers frequently offer more than the lowest possible amount above zero. The second dimension concerns *fairness*. Fairness is defined as making an offer that one would oneself be willing to accept, and requires the demand to be lower than one's own offer (*demand* < *offer*)[10]. Thus defined, fair strategies offer more than they themselves demand. Not only do they play well against themselves, but they also do not expect something that they are not willing to give[11,12]. Note that fairness differs from inequity aversion[26]. Individuals can have advantageous or disadvantageous inequity aversion which, in the extreme case, causes them either to split the resource equally or to expect equal shares, which in the UG would be similar to a strategy of *offer* = 0.5 and *demand* = 0.5. Fairness also differs from altruism, which would require an altruistic individual to perform actions that improve the fitness of another individual at its own expense. In fact, humans do not necessarily demand less than they offer themselves, and thus do not play *fairly* as defined by Rand *et al.*[10]. Third and finally, it seems reasonable to assume that offers above 0.5 will nearly always be accepted[27], across cultural differences and with few exceptions[28,29]; thus, humans are *demanding* (*demand* > 0.5), in that they accept all offers above 0.5 and have a diminishing acceptance rate for offers below 0.5.

Results

In the UG players either make an offer or accept/reject an offer made by the other player, and it is these values (*offer* and *demand*) that we evolved in the computational model. In order to compare the results of the model with human play, we have to express human play as those two values. First, we asked how demanding humans really are. When observing humans, one can gauge a probability distribution for the offer (p_{offer}(*offer*)), whereas the demand is typically reported as an acceptance rate dependent on the offer ($P_{acceptance}$(*offer*)). Inferring demand from aggregated acceptance rates is problematic, however, because the acceptance rate is a function of the offers in question and thus not an independent measure.

For illustration, imagine an UG in which proposers make only unusually high offers between 0.75 and 0.85. Using the conventional approach of averaging over measured acceptance rates, which here would be 1.0 in each case, would result in a high demand (~0.8). Now assume an UG in which the same two responders encounter two proposers who make extremely low offers, say, 0.2 and 0.4. If we averaged across acceptance rates, the same two responders would suddenly have a much lower demand. As an alternative approach, we will show how to measure the demand of a responder *independently* of the offer distribution. When demand was measured in this way, we found that humans do not play fairly. Specifically, the demand was higher than the player's own offer, violating the aforementioned definition of fairness[10].

In the present analysis, we quantified demand by using data from an actual study of the UG but *not* conditioning it on offers[30] (the findings in the present control group were consistent with those of many other experiments[2,17–24]). Specifically, we define the demand as the threshold that accepts half of all offers and rejects the other half. We fitted the data from the study together with an acceptance rate of 1.0 for offers above 0.5, because those offers are nearly always accepted (see above), to the logistic equation:

$$p(offer) = \frac{1.0}{1.0 + \exp^{(-(k-(offer-x_o)))}},$$

(1)

resulting in an optimal fit for $k = 15.52$ (the slope of the logistic function) and $X_0 = 0.21$ (the x-value of the sigmoid's midpoint). (The dotted line f(*offer*) in Fig. 1A shows the fit.) The *demand* was then found by solving the following equation for *offer*:

$$\int_0^{demand} f(offer)\, d_{offer} = \frac{1}{2} \int_0^1 f(offer)\, d_{offer}.$$

(2)

In contrast to previous studies quantifying the demand given specific offers (usually below 0.5), we found that the demand was well above 0.5 (see Fig. 1A).
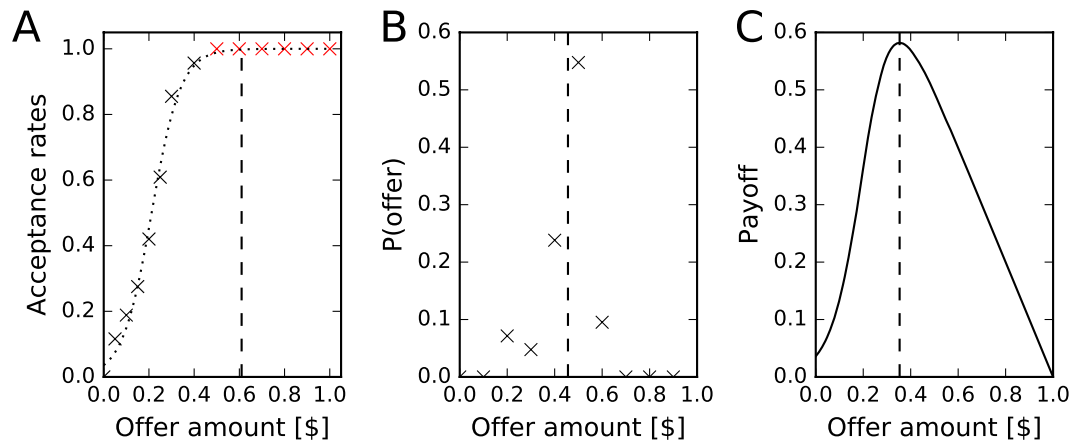
Figure 1. (**A**) Observed acceptance rates for offers between 0 and 0.5 in black, together with acceptance rates of 1.0 for offers above 0.5 in red (data from Harlé *et al.*[30]); the data were fitted to a logistic equation (Eq. 2, dotted line). The demand is indicated by the dashed line at 0.61. It indicates the threshold at which half of all possible offers, described by the fitted logistic function, were accepted. (**B**) Offer distribution from Forsythe *et al.*[31] with a mean of 0.456 (dotted line), indicating that humans offered on average about 0.45. (**C**) Dependence of payoff (y axis) given a particular *offer* (x axis) defined by Equation 3, using the demand distribution *f*(*offer*) (shown in A). This curve shows that an offer of around 0.38 (indicated by the dashed line) had the highest expected return given the demand distribution estimated in A.

Having described the demand of humans as a single threshold, we next asked whether this threshold was optimal given the measured offer distribution. Offers are normally distributed around a mean somewhere between 0.3 and 0.5 (depending on the study[10,31]; the data shown in Fig. 1B are from Forsythe *et al.*[31]). Assuming the demand distribution we inferred above, which included offers above 0.5, we assessed which offer would return the highest payoff. When *f*(*offer*) defines the likelihood that an offer will be rejected, the payoff can be quantified as:

$$payoff_{offer} = (1 - offer)f(offer).$$

(3)

We found the payoff to be maximal at an offer of $\approx$0.38 (see Fig. 1C). Thus, the mean offer made by humans turned out to be optimal for the demand we measured. Therefore, we suggest that the strategy humans play is roughly (0.4, 0.6). This strategy is far from the rational Nash equilibrium (*offer* > 0), inasmuch as it is *generous* (*offer* $\gg$ 0) as well as *demanding* (*demand* > 0.5) and *unfair* (*offer* < *demand*). However, if played against an opponent who meets the demands and accepts the offer, this strategy provides an advantage. Specifically, it offers less than half of the resource, demands more than half of it, and results in a payoff that is larger than that of the opponent. Interestingly, this finding also follows our intuition about what it takes to win this game: As the proposer, one should hand over less than one retains (*offer* < 0.5); as the responder, one should receive more than the proposer retains (*demand* > 0.5).

Of course, the strategy proposed here is dependent on the context of the game it is played in ref. 27. For example, the exact probabilities can depend on gender[20] as well as on hormonal levels[23,32]. However, the behaviour observed in these cases deviated even more from the rational Nash equilibrium than the proposed strategy. Our analysis was conducted to approximate the strategy humans play, and taking other contextual effects into account would not change our conclusions.

Generous strategies are more adaptive. What drives humans to deviate from the rational Nash equilibrium? Social preference models explain the deviation by assuming that humans are motivated by considerations other than purely self-regarding ones, such as inequity aversion[33,34], reciprocity[35,36], fairness[10,17,18,34,37], altruism, and punishment[38–43]. Alternatively, it has been argued that few real-world games are truly anonymous and one-shot[44–47], or that during evolution games are repeated over generations, even though individuals might be confronted with "once in a lifetime" decisions (for details, see Binmore 2012[48]). Repeated encounters allow players to deploy different strategies, so that, for example, players stop cooperating with notorious defectors. Similarly, the green-beard effect[49,50] could affect the UG. The green-beard effect describes a situation in which a particular trait (the green beard) is recognized by other organisms, and permits conclusions to be drawn about organisms with that trait. In the UG, players recognized in this way might elicit different responses—for example, a particularly demanding strategy might receive a higher offer. The notion of repeated games is independent from kin detection, because the latter can also operate in one-shot games and independently of a green-beard effect. In sum, all these accounts help to explain why humans do not play the rational Nash equilibrium, and why their behaviour is inconsistent with the canonical economic model that considers humans to be rational and self-regarding.

Complementing these accounts, we propose that humans, being social animals, evolved behaviour that is highly influenced by kin selection[49,51]. Specifically, we suggest that humans assume that the anonymous opponent is not likely to be kin-related. Consequently, they play a strategy in the UG that is optimal against non-kin
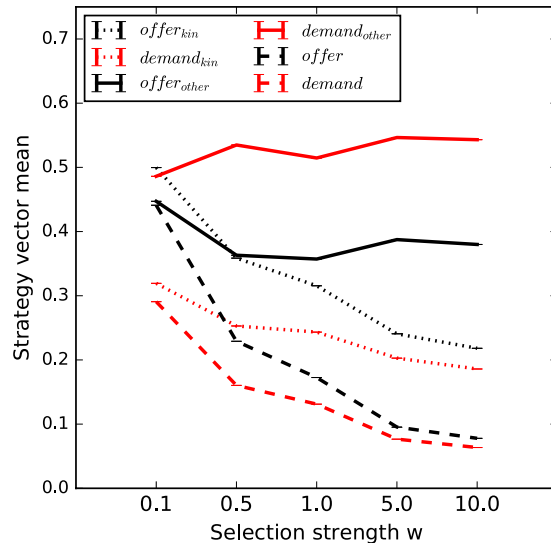
Figure 2. Mean evolved *offer* (black) and *demand* (red) at different selection strength (*w*, x axis), for two conditions. Dashed lines show strategies evolved without kin detection and resemble the findings from Rand *et al.*[10]. For strategies permitted to evolve kinship detection, dotted lines represent *offer*$_{kin}$ and *demand*$_{kin}$ when playing against kin, whereas full lines represent *offer*$_{other}$ and *demand*$_{other}$ when playing against others detected as non-kin. Error bars indicate standard error. Due to the large number of data points, error became vanishingly small.

opponents. Humans and their evolutionary ancestors were capable of detecting kin, and there is ample evidence for an evolved capacity to detect kinship relations in humans and animals[52]. Next, we will show that strategies that cannot discern between kin and non-kin evolve such that they play the rational Nash equilibrium. In contrast, strategies that are empowered to play differently against kin and non-kin, based on their ability to detect kinship relations[49,51], evolve to be *generous* (*offer* $\gg$ 0.5), *unfair* (*offer* < *demand*), and *demanding* (*demand* > 0.5).

Kin detection, kin selection. In evolving strategies to play the anonymous one-shot UG without the ability to detect kinship, we confirmed the results from Rand *et al.*[10] (see dashed lines in Fig. 2): Agents evolved the behaviour defined by the rational Nash equilibrium. Furthermore, increasing the strength of selection increased the propensity of agents to play the thus defined rational strategy. Consequently, the evolution of strategies unable to detect kinship cannot explain why humans play in a way that is generous, unfair, and demanding[10].

In the next step, we introduced the ability to detect kinship. In the computational evolutionary model, each agent received a *tag* $\in$ [0, 1] and a threshold to recognize kin $t \in$ [0, 1][53,54]. Let us imagine two strategies, *A* and *B*. Strategy *A* will recognize strategy *B* as kin if the difference between the two tags is smaller than the recognition threshold of *A* ($|tag_A - tag_B| < t_A$). Note that strategy *B* could come to a different conclusion than *A*. For illustration, assume $tag_A = 0.5$ and $tag_B = 0.7$, so that the difference between their tags is 0.2. If the threshold of *A* is $t_a = 0.3$, *A* will recognize *B* as kin. However, *B*'s threshold might be stricter ($t_B = 0.1$), meaning that *B* does not recognize *A* as kin. To permit strategies to play conditionally on the kin tag, we had strategies use one pair of values (*offer*$_{kin}$, *demand*$_{other}$) when recognizing the other party as kin and another pair of values (*offer*$_{other}$, *demand*$_{other}$) when recognizing the other party as non-kin. Note that if a strategy experienced a mutation of the *tag* in such a way that it became close enough to the *tag* of another strategy, it would recognize it as kin. In this way, non-kin strategies can be misinterpreted as kin.

When evolving strategies that are able to detect kin, we have to distinguish between strategies that classify the other party as kin and strategies that classify the other party as non-kin. If the evolved strategy decides that the other party is kin, its behaviour will be similar to that predicted by the rational Nash equilibrium (see dotted lines in Fig. 2). However, if the evolved strategy decides that the other party is non-kin, a much more generous strategy with a high demand will evolve (see solid lines in Fig. 2). Based on the latter decision, the evolved strategy will produce behaviour that is very similar to that observed in humans (see also Fig. 1), namely the play of this evolved strategy (0.4, 0.6) will be *generous* (*offer* $\gg$ 0.0), *unfair* (*offer* < *demand*), and *demanding* (*demand* > 0.5).

Although the *tag* did not converge on a particular value (data not shown), we found that for higher strength of selection (*w* > 0.1) the threshold (*t*) evolved to be small, indicating that agents were indeed selected to detect kin (see Fig. 3). The noise on the threshold was uniform and can be explained by random mutations[55], indicating that the threshold wasn't under strong. It seems likely that successful strategies gave rise to offspring with identical tags which in turn did not require low thresholds. Once the tag mutates in an individual, the threshold needs to be low again in order to discriminate between the new mutants. This back and forth might explain the limited selection strength.

The ability to discern between kin and non-kin permitted players to deploy two different strategies. This has profound implications for the evolutionary stability of strategies[56]. Simply put, a strategy has to fare better against itself (its own kin) than other strategies that play against it to become evolutionarily stable (defined by
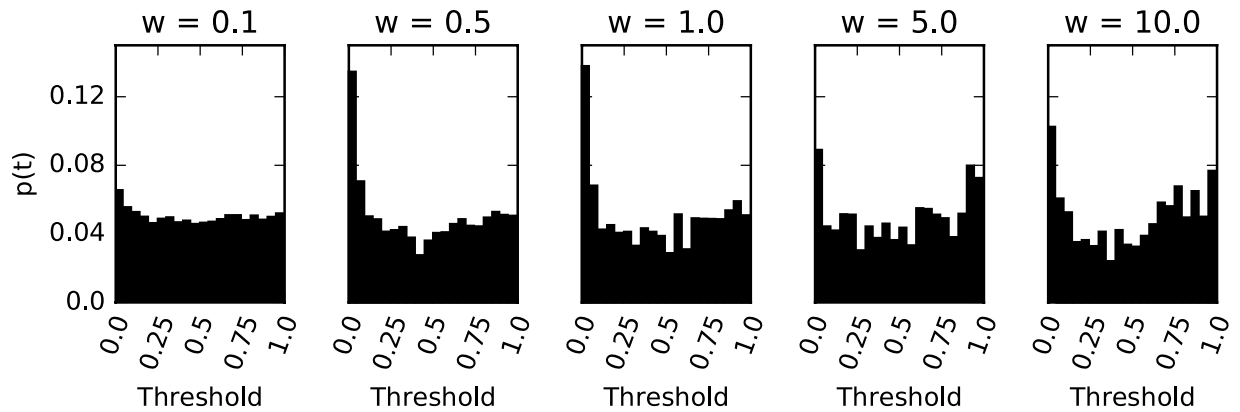
Figure 3. Distribution of evolved kin recognition thresholds. The histograms show the distribution of the kin recognition threshold recovered from the line of decent for different selection strength w ranging from 0.01 (left) to 10.0 (right) as annotated above. Except for the weakest selection strength we find a preference for low thresholds around 0.0.

Smith 1982[56], as the relation between payoffs E that strategies receive: $E(self, self) > E(other, self)$). By detecting kin, a strategy can maximize the payoff it receives against itself ($E(self, self)$), independently of what other strategies receive when playing it ($E(other, self)$). Therefore, evolved strategies in the UG maximized their payoff when playing against themselves by playing what the rational Nash equilibrium dictates. This behaviour is also robust against invasion by others who are misclassified as kin. At the same time, the branch of the strategy that played against non-kin had the ability to maximize its own payoff while minimizing the gains of others. This was achieved by offering less than 0.5 and demanding more than 0.5 (see the Supplementary Information for details). Increasing the minimal recognition threshold gradually compromised the ability to detect kin. Consequently, strategies cease to be generous, fair, and demanding. Ultimately, they returned to the strategy defined by the rational Nash equilibrium.

Discussion

Our results were obtained by computationally evolving strategies to play the UG conditionally on a kin tag and comparing the results with observations of human play. The model used a well-mixed population, and reduced behaviour in the UG to an offer and demand threshold. Interesting opportunities to expand the model in the future would be to include more complex strategies, iterated play, and spatial structures, but we do not expect any of these additions to fundamentally change our results. The model also assumed that humans use a single demand threshold, which we derived from experimental data for offers below 0.5, and we assumed all offers above 0.5 to be accepted. Future experiments that measure the demand independently from the actual offer (e.g., with the help of the strategy method[57]) might be able to model human behaviour more accurately, and would thus allow a more precise comparison between computationally evolved strategies and human behaviours.

We found that humans play the very same strategy in the UG that computationally evolved strategies played against non-kin others. We therefore suggest that humans, thrown into the UG, assume the anonymous other to be non-kin. In addition, we found that evolved strategies that assumed the other to be kin played a strategy similar to that described by the rational Nash equilibrium. Both of these findings are consistent with findings about the impact of genetic relatedness and with the offer and demand distribution[22] in the UG. For instance, genetically highly related villagers in Dominica[58] have been found to demand much less (0.152) than genetically unrelated Amazon Mechanical Turk workers[10] ($0.25 < offer < 0.5$). There may also be other reasons for humans in small communities to play less demanding; these possible reasons deserve to be explored in the future.

Although these results are quite astonishing and offer an explanation for human behaviour in the UG, the computational model is abstract. Of course, humans have not evolved to play the UG, and we cannot rule out other influences on human behaviour in this game. However, the computational model is general enough to be applicable to all organisms that have the ability to detect kin. Our analysis suggests that kin detection not only benefits kin, but also enables the evolution of generosity towards others who are not kin.

Our results also offer a new interpretation for social preference models. The robust finding that humans routinely deviate from the rational Nash equilibrium in the UG has led some to conclude that human behaviour is not fully rational in the terms of the canonical economic model. Social preference models have been proposed to explain the lack of full rationality. Here, however, we find that strategies that can distinguish between kin and non-kin evolve their *offer* and *demand* to be similar to those found in humans. This suggests that the human strategy is not necessarily irrational; rather, that is it the rational strategy if kinship relations are considered. Thus, social preference models might capture the psychology that enables humans to play the evolutionarily optimal strategy.

Methods

The one-off UG was implemented as described[10] before. Kinship detection was also added as described[53,54] before. Two-hundred independent evolutionary runs were performed per condition, using independent random number seeds for each replicate. The population was started with random genomes. Population size was 100 and the

mutation rate was 0.01 for *offer* and *demand* as well as for the tag and the recognition threshold. The mutation rate was implemented as a per site mutation rate, so that each value of the strategy vector could mutate independently. If a mutation occurred, a new value was drawn from a uniform random distribution [0, 1]. At the end of each run, the line of descent (LOD) was reconstructed[59] and the results were averaged over the LOD. Using the population mean instead of the LOD made no significant difference to the results (data not shown). The code and the analysis software can be downloaded from https://github.com/ahnt/UltimatumGame.

References

1. van Damme, E. *et al.* How Werner Güth's ultimatum game shaped our understanding of social behavior. *J. Econ. Behav. Organ.* **108,** 292–318 (2014).
2. Güth, W., Schmittberger, R. & Schwarze, B. An experimental analysis of ultimatum bargaining. *J. Econ. Behav. Organ.* **3,** 367–388 (1982).
3. Page, K. M., Nowak, M. A. & Sigmund, K. The spatial ultimatum game. *Proc. R. Soc. Lond. B Biol. Sci.* **267,** 2177–2182 (2000).
4. Szolnoki, A., Perc, M. & Szabó, G. Defense mechanisms of empathetic players in the spatial ultimatum game. *Phys. Rev. Lett.* **109,** 078701 (2012).
5. Kuperman, M. & Risau-Gusman, S. The effect of the topology on the spatial ultimatum game. *Eur. Phys. J. B* **62,** 233–238 (2008).
6. Sinatra, R. *et al.* The ultimatum game in complex networks. *J. Stat. Mech. Theor. Exp.* **2009,** P09012 (2009).
7. da Silva, R., Kellermann, G. A. & Lamb, L. C. Statistical fluctuations in population bargaining in the ultimatum game: static and evolutionary aspects. *J. Theor. Biol.* **258,** 208–218 (2009).
8. Ichinose, G. & Sayama, H. Evolution of fairness in the not quite ultimatum game. *Sci. Rep.* **4** (2014).
9. Mitzkewitz, M. & Nagel, R. Experimental results on ultimatum games with incomplete information. *Int. J. Game Theory* **22,** 171–198 (1993).
10. Rand, D. G., Tarnita, C. E., Ohtsuki, H. & Nowak, M. A. Evolution of fairness in the one-shot anonymous ultimatum game. *Proc. Natl. Acad. Sci. USA* **110,** 2581–2586 (2013).
11. Nowak, M. A., Page, K. M. & Sigmund, K. Fairness versus reason in the ultimatum game. *Science* **289,** 1773–1775 (2000).
12. Nowak, M. A. & Page, K. M. Empathy leads to fairness. *Bull. Math. Biol.* **64,** 1101–1116 (2002).
13. Nash, J. F. Equilibrium points in n-person games. *Proc. Natl. Acad. Sci. USA* **36,** 48–49 (1950).
14. Gale, J., Binmore, K. G. & Samuelson, L. Learning to be imperfect: the ultimatum game. *Games Econ. Behav.* **8,** 56–90 (1994).
15. Rubinstein, A. Perfect equilibrium in a bargaining model. *Econometrica* **50,** 97–109 (1982).
16. Güth, W. & Tietz, R. Ultimatum bargaining behavior: a survey and comparison of experimental results. *J. Econ. Psychol.* **11,** 417–449 (1990).
17. Kahneman, D., Knetsch, J. L. & Thaler, R. H. Fairness and the assumptions of economics. *J. Bus.* **59,** S285–S300 (1986).
18. Straub, P. G. & Murnighan, J. K. An experimental investigation of ultimatum games: information, fairness, expectations, and lowest acceptable offers. *J. Econ. Behav. Organ.* **27,** 345–364 (1995).
19. Larrick, R. P. & Blount, S. The claiming effect: why players are more generous in social dilemmas than in ultimatum games. *J. Pers. Soc. Psychol.* **72,** 810–825 (1997).
20. Solnick, S. J. Gender differences in the ultimatum game. *Econ. Inq.* **39,** 189–200 (2001).
21. Camerer, C. F. *Behavioral Game Theory: Experiments in Strategic Interaction* (Princeton Univ. Press, 2003).
22. Wallace, B., Cesarini, D., Lichtenstein, P. & Johannesson, M. Heritability of ultimatum game responder behavior. *Proc. Natl. Acad. Sci. USA* **104,** 15631–15634 (2007).
23. Zak, P. J., Stanton, A. A. & Ahmadi, S. Oxytocin increases generosity in humans. *PLoS ONE* **2,** e1128 (2007).
24. Spiegel, M., Currie, J., Sonnenschein, H. & Sen, A. Understanding when agents are fairmen or gamesmen. *Games Econ. Behav.* **7,** 104–115 (1994).
25. Fehr, E. & Schmidt, K. M. The economics of fairness, reciprocity and altruism-experimental evidence and new theories. In *Handbook of the Economics of Giving, Altruism and Reciprocity* Ch. 8, 615–691 (Elsevier, 2006).
26. Fehr, E. & Schmidt, K. M. A theory of fairness, competition, and cooperation. *Q. J. Econ.* **114,** 817–868 (1999).
27. Camerer, C. F. & Fehr, E. When does "economic man" dominate social behavior? *Science* **311,** 47–52 (2006).
28. Henrich, J. *et al.* In search of homo economicus: behavioral experiments in 15 small-scale societies. *Am. Econ. Rev.* **91,** 73–78 (2001).
29. Henrich, J. *et al. Foundations of Human Sociality-Economic Experiments and Ethnographic: Evidence From Fifteen Small-Scale Societies* (Oxford Univ. Press, 2004).
30. Harlé, K. M., Allen, J. J. B. & Sanfey, A. G. The impact of depression on social economic decision making. *J. Abnorm. Psychol.* **119,** 440–446 (2010).
31. Forsythe, R., Horowitz, J. L., Savin, N. E. & Sefton, M. Fairness in simple bargaining experiments. *J. Econ. Behav. Organ.* **6,** 347–369 (1991).
32. Zak, P. J. *et al.* Testosterone administration decreases generosity in the ultimatum game. *PLoS ONE* **4,** e8330 (2009).
33. Bolton, G. E. & Ockenfels, A. ERC: a theory of equity, reciprocity, and competition. *Am. Econ. Rev.* **90,** 166–193 (2000).
34. Fehr, E. & Schmidt, K. M. A theory of fairness, competition, and cooperation. *Q. J. Econ* **114,** 817–868 (1999).
35. Dufwenberg, M. & Kirchsteiger, G. A theory of sequential reciprocity. *Games Econ. Behav.* **47,** 268–298 (2004).
36. Rabin, M. Incorporating fairness into game theory and economics. *Am. Econ. Rev.* **83,** 1281–1302 (1993).
37. Van Dijk, E. & Vermunt, R. Strategy and fairness in social decision making: sometimes it pays to be powerless. *J. Exp. Soc. Psychol.* **36,** 1–25 (2000).
38. Davis, D. D. & Holt, C. A. *Experimental Economics* (Princeton Univ. Press, 1993).
39. Sigmund, K., Hauert, C. & Nowak, M. A. Reward and punishment. *Proc. Natl. Acad. Sci. USA* **98,** 10757–10762 (2001).
40. Fehr, E. & Gächter, S. Altruistic punishment in humans. *Nature* **415,** 137–140 (2002).
41. Boyd, R. The evolution of altruistic punishment. *Proc. Natl. Acad. Sci. USA* **100,** 3531–3535 (2003).
42. Hintze, A. & Adami, C. Punishment catalyzes the evolution of cooperation. E-print: arXiv:1210.5233 (2012).
43. Hintze, A. & Adami, C. Punishment in public goods games leads to meta-stable phase transitions and hysteresis. *Phys. Biol.* **12,** 046005 (2015).
44. Binmore, K. Economic man-or straw man? *Behav. Brain Sci.* **28,** 817–818 (2005).
45. Samuelson, L. Foundations of human sociality: a review essay. *J. Econ. Lit.* **43,** 488–497 (2005).
46. Binmore, K. G. *Does Game Theory Work? The Bargaining Challenge* (MIT Press, 2007).
47. Woodward, J. Social preferences in experimental economics. *Philos. Sci.* **75,** 646–657 (2008).
48. Binmore, K. *Evolution and Rationality* (Cambridge Univ. Pres, 2012).
49. Hamilton, W. D. The genetical evolution of social behaviour. II. *J. Theor. Biol.* **7,** 17–52 (1964).
50. Dawkins, R. *The Selfish Gene* (Oxford Univ. Press, 2006).
51. Hamilton, W. D. The genetical evolution of social behaviour. I. *J. Theor. Biol.* **7,** 1–16 (1964).
52. Lieberman, D., Tooby, J. & Cosmides, L. The architecture of human kin detection. *Nature* **445,** 727–731 (2007).
53. Riolo, R., Cohen, M. & Axelrod, R. Evolution of cooperation without reciprocity. *Nature* **414,** 441–443 (2001).

54. Traulsen, A. & Schuster, H. G. Minimal model for tag-based cooperation. *Phys. Rev. E. Stat. Nonlin. Soft Matter Phys.* **68,** 046129 (2003).
55. Zeeman, E. C. Dynamics of the evolution of animal conflicts. *Journal of theoretical Biology* **89,** 249–270 (1981).
56. Smith, J. M. *Evolution and the Theory of Games* (Cambridge Univ. Press, 1982).
57. Selten, R. Die strategiemethode zur erforschung des eingeschränkt rationalen verhaltens im rahmen eines oligopolexperiments. In *Beiträge zur experimentellen Wirtschaftsforschung,* 136–168 (JCB Mohr, 1967).
58. Macfarlan, S. & Quinlan, R. Kinship, family, and gender effects in the ultimatum game. *Hum. Nat.* **19,** 294–309 (2008).
59. Lenski, R. E., Ofria, C., Pennock, R. T. & Adami, C. The evolutionary origin of complex features. *Nature* **423,** 139–144 (2003).

Acknowledgements

Author Contributions

A.H. and R.H. designed the experiments, interpreted the results, and wrote the manuscript. A.H. wrote all code and performed all computational experiments and analysis.

Additional Information

Supplementary information accompanies this paper at http://www.nature.com/srep

Competing financial interests: The authors declare no competing financial interests.

How to cite this article: Hintze, A. and Hertwig, R. The Evolution of Generosity in the Ultimatum Game. *Sci. Rep.* **6**, 34102; doi: 10.1038/srep34102 (2016).