Evolution of fairness in the one-shot anonymous Ultimatum Game

David G. Rand[a,b,c,1], Corina E. Tarnita[a,d,e,1], Hisashi Ohtsuki[f,g], and Martin A. Nowak[a,h,i,2]

[a]Program for Evolutionary Dynamics, [b]Department of Psychology, [d]Society of Fellows, and Departments of [h]Organismic and Evolutionary Biology and [i]Mathematics, Harvard University, Cambridge, MA 02139; [c]Department of Psychology, Yale University, New Haven, CT 06520; [e]Department of Ecology and Evolutionary Biology, Princeton University, Princeton, NJ 08544; [f]Department of Evolutionary Studies of Biosystems, School of Advanced Sciences, The Graduate University for Advanced Studies, Hayama, Kanagawa 240-0193, Japan; and [g]Precursory Research for Embryonic Science and Technology, Japan Science and Technology Agency, Kawaguchi, Saitama 332-0012, Japan

Classical economic models assume that people are fully rational and selfish, while experiments often point to different conclusions. A canonical example is the Ultimatum Game: one player proposes a division of a sum of money between herself and a second player, who either accepts or rejects. Based on rational self-interest, responders should accept any nonzero offer and proposers should offer the smallest possible amount. Traditional, deterministic models of evolutionary game theory agree: in the one-shot anonymous Ultimatum Game, natural selection favors low offers and demands. Experiments instead show a preference for fairness: often responders reject low offers and proposers make higher offers than needed to avoid rejection. Here we show that using stochastic evolutionary game theory, where agents make mistakes when judging the payoffs and strategies of others, natural selection favors fairness. Across a range of parameters, the average strategy matches the observed behavior: proposers offer between 30% and 50%, and responders demand between 25% and 40%. Rejecting low offers increases relative payoff in pairwise competition between two strategies and is favored when selection is sufficiently weak. Offering more than you demand increases payoff when many strategies are present simultaneously and is favored when mutation is sufficiently high. We also perform a behavioral experiment and find empirical support for these theoretical findings: uncertainty about the success of others is associated with higher demands and offers; and inconsistency in the behavior of others is associated with higher offers but not predictive of demands. In an uncertain world, fairness finishes first.

cooperation | prosociality | stochastic dynamics

Game theorists traditionally assume that people act fully rationally to maximize their own financial gains. A wealth of behavioral data, however, has demonstrated that many people are influenced by the payoffs of others, exhibiting so-called "other-regarding preferences" (1). The Ultimatum Game (UG) has been a particularly influential example of this phenomenon (2). In the UG, two players have to divide a certain sum of money between them. One player (the proposer) makes an offer. The other player (the responder) can either accept the offer, in which case each receives the money as proposed, or reject the offer, in which case neither player receives anything. In a one-shot anonymous UG, a rational self-interested proposer will offer the minimum amount that she believes will be acceptable to the responder. A rational self-interested responder will accept any nonzero offer. Thus, under common knowledge of the rationality of both players, the subgame perfect Nash equilibrium is for the proposer to make the minimum possible offer, and for the responder to accept it (2).

To evaluate these predictions, many behavioral experiments have been conducted using the UG (1–8). Although there is considerable quantitative variation across studies, two clear qualitative deviations from rational self-interest are robustly observed: (i) many responders choose to reject low (but nonzero) offers, and (ii) many proposers offer more than the minimum amount required to avoid rejection. One popular explanation of both of

these findings is that people are motivated by a sense of fairness (or "inequity aversion"): Subjects prefer both players to receive equal payoffs, and are willing to pay a price to create more equitable outcomes (9). By this argument, responders who reject low offers incur a cost to avoid getting a smaller payoff than the proposer (disadvantageous inequity), and proposers who offer more than needed to avoid rejection incur a cost to avoid receiving a larger payoff than the responder (advantageous inequity). Additional evidence of this psychological principle is demonstrated by an experiment where subjects will pay to alter randomly assigned payoffs of others to induce greater equality (10).

Furthermore, it is typically observed that people are more averse to disadvantageous inequity than they are to advantageous inequity (9), and research with children finds that disadvantageous inequity develops earlier than advantageous inequity (11–14). These results suggest that the two forms of fairness are most likely cognitively distinct. Some have argued that proposer behavior can be entirely explained by strategic motivations: given that many responders reject low offers, it may be payoff maximizing to offer even splits (15, 16). Others, however, contend that fairness concerns play at least some part in the high offers of proposers (17, 18); a comprehensive review (1) concludes that high proposer offers are likely the result of a combination of strategic and fairness-based motivations.

Fairness presents a proximate psychological motivation for the observed behavior. What, however, is the ultimate evolutionary explanation for why we should have come to possess such fairness preferences? To explore the origins of fairness, we study an evolutionary process in which strategies with higher payoffs tend to become more common in the population (19–22). This process could describe genetic evolution, or cultural evolution through social learning, both of which have been linked to play in the UG (4, 18, 23). In the context of genetic evolution, agents reproduce and die, and mutations introduce variation into the gene pool. In the context of cultural evolution, individuals sometimes change strategy and copy the strategy of another, with higher payoff strategies being more likely to be imitated. Here mutation represents either experimentation, in which individuals try new random strategies, or confusion regarding the strategies used by other players (leading the imitator to adopt a strategy different from that of the imitatee) (24). This type of process based on reproduction is distinct from strategy update rules that use prospective reasoning, such as best response dynamics (25)

EVOLUTION

SOCIAL SCIENCES

(for a comparison of evolutionary dynamics with prospective reasoning, see ref. 26).

Using this evolutionary framework, we can explore the conditions under which natural selection leads to fair behavior. The classical approach to evolutionary game theory (27–29) is deterministic. Such models assume that higher payoff strategies always become more common, whereas lower payoff strategies always die out. Under deterministic game dynamics, evolution favors self-interest, and in the UG selection leads to the rational self-interested strategy where agents offer and demand nothing (30). To explain fairness using deterministic dynamics, it is therefore necessary to invoke some additional evolutionary mechanism.

One approach involves reputation formation (31–36). It has been shown that fairness can be favored by natural selection if agents can recognize their partners' strategies (37, 38) or have reputations that carry from game to game (30). Here it pays to reject a stingy offer today so that others will make you higher offers in the future. Without a sufficiently high expectation of future interactions and a sufficiently strong reputation system, however, fairness collapses. An alternate approach studies one-shot anonymous games but assumes an asymmetric mutation structure, such that proposers experiment with new strategies less often than responders: the greater variation in responder behavior forces proposers to make higher offers (39). A third deterministic approach involves one-shot anonymous games played among very small groups. Here payoff relative to your coplayers is critical, and so accepting unfair offers can put you at a disadvantage. Thus, it can be advantageous to reject unfair offers. The optimal demand, however, is inversely proportional to the number of coplayers, and thus is negligibly greater than zero in all but the smallest groups (40). Note that theories related to multilevel selection (41, 42) do not help explain fairness in the UG, as a group of individuals offering and accepting minimal offers receives the same average payoff as a group of "fair" players with nonzero offers and demands.

An important element which is not included in these previous analyses is that randomness plays a key role in the course of evolution, especially in finite populations. Agents might be involved in many different games, such that their payoff in the UG contributes only a small amount to their total fitness (43). Alternatively, individuals may make errors in social learning, due to issues such as bounded rationality (44) and difficulties in correctly assessing others' payoffs. In either situation, lower payoff strategies may sometimes spread through the population by chance despite their relative disadvantage, and higher payoff strategies may die out. Such stochastic effects can have potentially dramatic effects on evolutionary outcomes (26, 43, 45). In the present paper, we explore the evolution of strategies in the UG in finite populations, studying the whole spectrum of selection intensities ranging from the limit of weak selection (where reproduction is almost completely random) to strong selection (where higher payoff strategies almost always increase in frequency). We show that when selection is not too strong, evolution can lead to the nonzero rejections and generous offers observed experimentally, without the need for any additional evolutionary mechanisms such as reputation systems, and with no a priori assumptions about asymmetries or other-regarding preferences. Self-interested natural selection in finite populations favors the evolution of fairness when sufficient randomness is present.

We model the UG by imagining two players who have to split an amount summing to unity. In any given interaction, players are randomly assigned to the roles of proposer and responder. We specify an agent's strategy with two parameters p and $q \in [0,1]$, where p is the amount offered when acting as proposer, and q is the minimum amount demanded when acting as responder, or the "rejection threshold." An offer p is accepted by a responder with the minimum demand q if and only if $p \geq q$.

Therefore, the average payoff for a player using strategy (p_1,q_1) interacting with a player using strategy (p_2,q_2) is given (up to the factor 1/2, which we henceforth omit) by (i) $1-p_1 + p_2$, if $p_1 \geq q_2$ and $p_2 \geq q_1$; (ii) $1-p_1$, if $p_1 \geq q_2$ and $p_2 < q_1$; (iii) p_2, if $p_1 < q_2$ and $p_2 \geq q_1$; and (iv) 0, if $p_1 < q_2$ and $p_2 < q_1$.

We consider the stochastic evolution of strategies in a population of finite size N. Each player i plays the UG with each of the $N-1$ other players, and receives an average payoff π_i. Player i's effective payoff (or fecundity) is then defined as $\exp[w\pi_i]$, where w is called the "intensity of selection." An intuition behind this effective payoff function is that the higher the intensity of selection, the more likely agents with higher payoffs are to be imitated (to reproduce). At the extreme of $w \to \infty$, only those who obtain the highest payoff are imitated (strong selection). The other extreme $w \to 0$ is called the weak selection limit; in this case, all strategies have almost the same effective payoff and the dynamics is dominated by neutral drift. We study the Moran process (46, 47), where in each generation an individual is randomly picked to change strategy (die), and another individual is picked proportional to effective payoff to be imitated (reproduce). With probability u, a mutation occurs and instead a random strategy is chosen. We begin by considering global mutation, in which a mutant's p and q are independently drawn from the uniform distribution in [0, 1], and later show that the results are qualitatively unchanged when instead we use local mutation.

The dynamics depends significantly on the mutation rate u. In the low mutation limit $u \to 0$, a novel mutant will either die out or completely take over the population before a new mutant arises (43, 48–51). Thus, the population transitions between homogeneous states, in which all agents in the population play the same strategy at any given time. Here, strategies which can protect themselves from invasion do best. Conversely, in the high mutation limit $u \to 1$, all strategies are present at approximately equal abundances at the same time (24, 52). Thus, success is determined not by resisting invasion, but by performing best when playing against all strategies with equal probability (i.e., playing a randomly selected opponent). Intermediate mutation rates result in intermediate outcomes between these two dynamical extremes (for technical details, see Methods and Supporting Information).

Results

We begin with agent-based simulations using population size $N = 100$, and vary the intensity of selection w as well as the mutation rate u. For each set of simulation parameters we determine the steady-state frequency distribution over the [p, q] space. First we ask which strategy is favored by natural selection (i.e., is most common in the population). We find that the most common strategy displays $p > q > 0$ as long as selection strength w is not too large (Fig. 1). Thus, selection favors both aspects of fairness observed in behavioral experiments: responders make nonzero demands (disadvantageous inequity aversion, $q > 0$) and proposers offer more than is necessary to avoid rejection (advantageous inequity aversion, $p > q$). Evolution in finite populations can select for fairness, without needing to invoke any additional mechanisms.

We now turn from the most common strategy to consider the average (mean) strategy. Fig. 2 shows how the time-averaged values of p and q vary systematically with changes in selection strength w and mutation rate u. When selection is very weak, the dynamics is dominated by neutral drift, and mean p and q are both ≈ 0.5. As the selection intensity increases, both p and q decrease, approaching the rational self-interested strategy $p = q = 0$. Critically, however, q decreases faster than p. Thus, in Fig. 2 we observe both aspects of fair behavior, $p > q$ and $q > 0$, across a wide range of parameter values. As with the modal strategy considered in Fig. 1, we see in Fig. 2 that although the selection strength and mutation rate quantitatively affect the mean p and q,

Fig. 1. With intermediate selection and mutation, the most common strategy is fair, having $p > q > 0$. Shown are the frequencies of $[p,q]$ pairs averaged over 10^8 generations. To aid convergence, the p and q values of agents in the simulations in the figure are discretized in increments of 1/12 (all other simulations use a continuous strategy space). Red indicates high frequency; yellow indicates lower frequency. The most common strategy is indicated with a black ×. Simulations use $n = 100$ and $u = 10^{-1.25}$, with w varying across $10^{-1.5}$ (A), 10^{-1} (B), $10^{-0.5}$ (C), and 10^2 (D). Similar results are obtained using other mutation rates (Fig. S1). Note that strong selection drives the population to the smallest possible nonzero value of $p = q = 1/12$ in D (rather than $p = q = 0$), for the following reason: although $p = q = 1/12$ is neutral with $p = 1/12$, $q = 0$, the latter strategy can be invaded by $p = q = 0$, which in turn is risk-dominated by $p = 1/12$, $q = 1/12$; due to the discretized strategy space, no intermediate strategies exist.

the qualitative result is general (outside of the weak and strong selection extremes): The mean proposal p is greater than the demand q, and the average demand q is substantially greater than 0. These findings are very different from the results of classical evolutionary game theory using either the replicator equation studying the mini UG (30) or adaptive dynamics studying the full UG (53), where the population converges to the rational self-interested strategy $p = q = 0$ unless other mechanisms are present.

Our results are robust to the manner in which mutants are selected. Replacing the uniform mutation described above with a local


Fig. 2. Across a wide range of selection strengths and mutation rates, evolution results in fairness on average: the mean minimum amount demanded has $q > 0$, and the mean offer has $p > q$. Shown are time-averaged values of p and q over 10^8 generations, using the population size $n = 100$ and mutation rate (A) $u = 10^{-3}$, (B) $u = 10^{-2}$, and (C) $u = 10^{-1}$. Shown in yellow are the parameter regions which agree with experimental data, $0.3 < p < 0.5$ and $0.2 < q < 0.35$, based on additional simulations examining selection strengths in increments of 0.1.

mutation scheme in which mutants are some random perturbation from the parent strategy gives qualitatively similar results. See SI Local Mutation Kernel and Figs. S2 and S3 for details.

A pessimistic interpretation of the results in Fig. 2 is as follows. Perhaps selection always favors the rational self-interested strategy $p = q = 0$, and the fact that the mean p and q transition from 0 to 0.5 as w decreases is the trivial result of increasing neutral drift driving the mean away from the optimal (selfish) strategy. On the contrary, however, we clearly see that this is not the case. Instead, the frequency distributions in Fig. 1 are centered around fair strategies with large offers and demands, as long as selection is not too strong. Put differently, the strategy most favored by natural selection is the strategy that is most common under mutation-selection balance; thus, the fact that $p = q = 0$ is not the most common strategy when selection is not so strong shows that we truly are observing natural selection favoring fairness.

Thus far we have shown that evolution in finite populations can qualitatively reproduce both the negative and positive aspects of fair behavior demonstrated in experiments. Now we ask whether there can also be quantitative agreement between our model and the range of behaviors observed in the experimental data. On the negative side of fairness (disadvantageous inequity aversion), whereas Homo economicus would accept any nonzero offer, the mean demand q is substantially greater than zero across experiments, tending to lie in the range $0.2 < q < 0.35$. On the positive side of fairness (advantageous inequity aversion), subjects also offer more than is demanded: across experiments, the average offer p is substantially higher than the average demand, typically in the range $0.3 < p < 0.5$ (see Fig. S4 for mean p and q values from numerous experiments).

We now compare these experimental data with the average values of p and q from our agent-based simulations, across a range of selection strengths and mutation rates. We see that the evolutionary outcomes for a number of parameter combinations are quantitatively consistent with the experimental data, having mean offers $0.3 < p < 0.5$ and mean demands $0.2 < q < 0.35$. The parameter regions which lead to this agreement are highlighted in yellow in Fig. 2. We see that increasing the mutation rate leads to a corresponding increase in the selection strength needed to reproduce the experimental behavior. This

balancing is required to conserve the level of randomness in the system, which is increased by higher mutation and decreased by stronger selection (the opposite is true for the relationship between selection strength and population size, as shown in Fig. S5). We see that with the correct level of randomness, our evolutionary simulations can quantitatively reproduce the range of average behavior observed in experiments. This agreement stands in contrast with classical economic approaches as well as deterministic evolutionary dynamics, and demonstrates the potential power of finite population evolutionary analysis for understanding human behavior.

In addition to average behavior, it is also of interest to consider how the distribution of individual-level behaviors shown in Fig. 1 compares with experimental data (see Fig. S6 for histograms of p and q separately, rather than the joint $[p,q]$ distribution shown in Fig. 1). We begin with proposer behavior. Our model produces a unimodal distribution of p values that drops off sharply when p rises above 0.5. This result is generally consistent with the findings of behavioral experiments, with the exception of the model having substantially more variation in offers than is typically seen in experiments (i.e., a wider distribution), and including a low but nonzero density of probability weight for offers above 0.5 (whereas virtually no subjects offer more than 0.5 in most experiments) (1). Turning to responder behavior, our model again produces a fairly broad unimodal distribution with relatively little probability weight above 0.5. It is harder to compare these results with experimental data as few studies provide distributions of minimum acceptable offers, and the few that do are not consistent with each other: both unimodal distributions (2) and bimodal distributions with modes at 0 and 0.5 (3) have been observed. Further exploration of individual-level behavior, both theoretically and experimentally, is an important direction for future study.

Discussion

To gain an intuition for the evolutionary success of fairness in our agent based simulations, we turn to mathematical calculations. In the weak selection limit, where the average abundance of all strategies is approximately equal (steady-state p and q are uniformly distributed on the unit square), it is possible to analytically determine which strategy is most common. We find that the most frequent strategy depends strongly on the rate at which mutations arise in the population (see SI Intuition on the Role of Mutation and Table S1).

In the high mutation limit, all strategies are present in the population simultaneously with approximately equal frequency. Hence, the optimum strategy is the one that maximizes its expected absolute payoff against a randomly chosen opposing strategy. As has been shown previously (30), it is intuitive that this strategy is (1/2, 0). The offer of $p = 1/2$ maximizes the proposer's expected payoff of $p(1-p)$ when playing against a randomly chosen opponent; and the demand $q = 0$ maximizes expected payoff as responder because any nonzero demand results in lost profit. Thus, in the high mutation limit natural selection favors the first aspect of fair behavior (advantageous inequity aversion), with proposal p greater than demand q.

In the low mutation limit, on the other hand, the population dynamics is very different. A new mutant will either die out or take over the resident population before another mutant arises. Thus, although all strategies are still present at equal frequency in the steady-state distribution when in the limit of weak selection, at most two strategies are ever present in the population at the same time. Therefore, what matters in the low mutation limit is resisting invasion by a single (randomly chosen) other strategy: it is not expected absolute payoff that determines success, but rather expected relative payoff in pairwise competition with a random opponent (52).

What strategy then maximizes expected relative payoff? One can see that $p = q$ is a logical first requirement for success in the low mutation limit: any strategy that offers more than it demands ($p > q$) loses in pairwise competition with mutant strategies which offer p_m such that $p > p_m > q$ (the p_m mutant is less generous than the resident, but still has its offers accepted); and any strategy that demands more than it offers ($p < q$) always rejects its own offers and is outcompeted by all mutant strategies with $p_m > q_m$. Thus, we focus our attention on the self-consistent set of strategies with $p = q$. Consider the interaction between a relatively fair strategy S_F with $p = q = x$, and a relatively unfair strategy S_U with $p = q = x-«$. Both strategies receive the full payoff of 1 when playing against themselves; but, when S_F meets S_U, surprisingly the more fair strategy receives the higher payoff (provided it offers less than half). When S_F is the proposer, her offer is accepted, and she earns $1-x$ whereas S_U earns x; when S_U is the proposer, her lower offer is rejected, and neither player earns anything. Thus, fairer strategies always earn more than less fair strategies when they interact pairwise as long as $x < 0.5$. However, when considering expected relative payoff against a random opponent, there is a tradeoff: the more you offer (up to 0.5), the more strategies you outearn, but the smaller your margin of success is in each pairing. This creates two opposing forces resulting from increasing your offer: the decreasing marginal payoff versus the increasing number of strategies you outperform. These two forces balance out at some intermediate, optimal value of x. We find that this balance is achieved at $p = q = 1/3$, and that this result continues to hold when lifting the restriction $p = q$. Thus, when mutations are rare, it pays to reject nonzero offers; in this case, the most common strategy has the second experimentally observed aspect of fairness (disadvantageous inequity aversion), with $q > 0$.

At intermediate mutation rates, the evolutionary dynamic has characteristics of both the fully heterogeneous and fully homogeneous extremes. Because of the somewhat heterogeneous nature of the population, $p > q$ is favored by selection; and due to the somewhat homogeneous nature of the population, $q > 0$ is also favored. We therefore find that for $Nu > 1$, the most frequent strategy has $p = (1 + Nu)/(4 + 2Nu)$ and $q = 1/(2 + Nu)$. As shown in Fig. 3, we see evolution favoring both qualitative attributes of experimentally observed human behavior that were so challenging for classical game theory's "economic man": nonzero rejection thresholds, $q > 0$, and proposer's generosity beyond what is necessary to avoid rejection, $p > q$. The same results hold if instead of one population in which any individual can be either a proposer or a responder in any given game, we consider two separate interacting populations, one of proposers and one of responders. See SI Single–Population Formulation for


Fig. 3. In the weak selection limit, the modal strategy is fair for intermediate mutation. Shown are the most common strategies p (blue) and q (red) as functions of the mutation rate, calculated analytically in the weak selection limit (see SI Single-Population Formation for details). For low mutation $0 \leq Nu \leq 1$, $p = q = 1/3$ is the most common strategy. As mutation increases, the optimum proposal increases to 1/2 and the optimum threshold decreases to 0. For intermediate mutation rates, we observe both key features of real-world ultimate game behavior: p > q > 0.

the weak selection limit analysis and SI Two–Population Formulation for the treatment of separate populations.

To summarize, we find that (i) weakening selection increases the favored demand q (and therefore also the favored offer p), and (ii) increasing the mutation rate increases the favored offer p but not demand q. We now use these theoretical findings to generate two testable predictions, and evaluate these predictions by running a behavioral experiment.

The first prediction stems from the result regarding weakening selection. In our model, weaker selection means that agents have a harder time assessing which others have the highest payoffs when choosing whom to imitate. Therefore, we would predict that people who developed their strategies in settings where it was more difficult to assess the successfulness of others would make both larger offers and larger demands.

The second prediction comes from the result regarding increasing mutation. In our model, higher mutation means that agents are more likely to change their strategy at random. Therefore, we would predict that people who developed their strategies in settings where the behavior of others is less consistent would make higher offers, but not higher demands.

To evaluate these two predictions, we conduct an experiment using the online labor market Amazon Mechanical Turk (54–61). We recruit n = 140 subjects from around the world to play a one-shot anonymous UG. In addition, these subjects are asked "Among those you interact with in daily life, how clear is it which people are more or less successful?" as a measure of the intensity of selection under which they developed their strategy, and "How accurate do you think first impressions are when judging other people?" as a measure of the consistency of others (i.e., the inverse of the mutation/experimentation rate) under which they developed their strategy.

The results validate both of our theoretical predictions. Fig. 4A shows that subjects who report less clarity about the successfulness of others offer more and demand more. Fig. 4B shows that subjects who report less consistency of others offer more but do not demand more. These results are confirmed by statistical analysis using linear regression with robust SEs, including appropriate controls. See Methods for further details.

We have shown that in finite populations, where dynamics are stochastic and evolutionary trajectories are influenced by chance, natural selection favors fairness in the one-shot anonymous UG.

Furthermore, we have validated the predictions of our analysis using behavioral experiments. Previous analyses of the UG have focused on situations where selection is strong and higher payoff strategies always produce more offspring (30, 39, 40). In these settings where evolution is deterministic, fairness needs additional mechanisms to evolve. When the role of random chance is included, however, the results are very different. Without any reputation systems (30, 37), asymmetries between proposers and responders (39), or a priori assumptions about other-regarding preferences (9), the self-interested process of evolution can lead to behaviors which defy classic rational self-interest models (26, 45). These stochastic evolutionary simulations for finite populations can quantitatively reproduce the range of human behaviors observed in the laboratory. [This is as true for cross-cultural results on UG play from small-scale societies as it is for play by Western undergraduates, with the exception of the few societies in which offers above 50% are consistently rejected (23).] Thus, we do not necessarily need to invoke additional mechanisms to provide an evolutionary account of the origins of fairness, as long as the system is not too deterministic.

The emotions, intuitions, and preferences which guide us toward generosity and righteous anger in the one-shot anonymous UG (62–65) may be the proximate biological (4) and/or cultural (23) implementations of behaviors which are advantageous in the presence of weak selection and imperfect learning (for further discussion of the role of intuition in economic games, see ref. 66). Thus, stochastic evolutionary dynamics may offer an explanation for why we have come to have such a preference for fairness. When populations are finite and selection is not too strong, evolution can be fickle: Fitter strategies sometimes die out, and less-fit strategies sometimes triumph. However, in this unfair world, myopic self-interest is vanquished whereas fairness triumphs.

Methods

Agent-Based Simulations. Our main results are produced using agent-based simulations. In our simulations, agents interact in a well-mixed population of constant size 100. Each agent i has a strategy vector $[p_i, q_i]$ specifying that agent's behavior when acting as proposer (offers p_i) and responder (demands q_i).

Each generation, every agent plays the UG with every other agent, once in the proposer role and once in the responder role, and the resulting payoff π_i is the average of the payoffs over all 99 pairings.

Then one agent is picked proportional to $\exp[w\pi]$ to reproduce, where w is the intensity of selection; and one agent is picked at random to die. With probability $1 - u$, the dead agent's strategy is replaced with the reproducing agent's strategy; with probability u, a mutation occurs and instead the dead agent's strategy is replaced with a randomly selected strategy. Thus, u is the mutation rate.

Each agent's strategy is initialized randomly at the beginning of the simulation, and the strategies of all agents are recorded over 10^8 generations.

For details of the weak selection analytical calculations, see SI Single-Population Formulation.

Behavioral Experiments. To assess our theoretical predictions, we conduct a behavioral experiment. We recruit n = 140 subjects using the online labor market Amazon Mechanical Turk (AMT; for an overview of running experiments on AMT, as well as a discussion of the value of combining behavioral experiments and theoretical models, see ref. 55). Commensurate with standard wages on AMT, subjects receive a $0.20 baseline payment for participating, and then play a UG in which $0.40 is at stake. For evidence that these low stakes do not compromise the validity of behavioral data, see ref. 57. Subjects read a set of instructions explaining the game, and are told they will be randomly assigned to be either the proposer or the responder. They are then asked to calculate the payoff received by the proposer and responder in two different scenarios to ensure that they understand the payoff structure. Only subjects who answer correctly are allowed to participate.

After clearing the comprehension questions, subjects indicate the minimum offer they would accept if they are assigned to be the responder. Then they indicate the amount they would offer if they are assigned to be the proposer. Finally, they complete a demographic questionnaire that includes the questions "Among those you interact with in daily life, how clear is it which people are more or less successful?" and "How accurate do you think
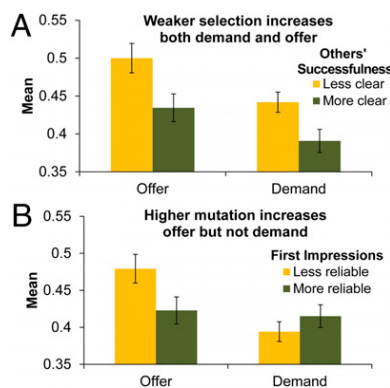

Fig. 4. A behavioral experiment in which subjects play a one-shot anonymous UG confirms two predictions of our model. (A) Subjects that report a less-clear understanding of who in their community is more versus less successful (i.e., that developed their strategies under weaker selection) make larger offers and larger demands. (B) Subjects that report first impressions to be less reliable (i.e., that developed their strategies under higher mutation rates) make larger offers but not larger demands. Error bars indicate SEM. For visualization, subjects are divided into two groups in each panel using a median split on question responses.

rst impressions are when judging other people?," each reported using a 5-point Likert scale (1 = Very unclear to 5 = Very clear for the rst question; 1 = Very inaccurate to 5 = Very accurate for the second).

Once all subjects have been recruited, they are randomly paired and assigned roles, the resulting payoffs are calculated, and each subject is paid accordingly using the AMT payment system. No deception is used. The practice of having subjects specify a strategy which dictates a decision in each possible outcome and then having actual payoffs determined by ex post matching, referred to as the "strategy method," is a common technique in experimental economics (and is used by all of the experimental papers whose data we visualize in Fig. S4). This is particularly true for eliciting responder behavior in the UG, as low proposer offers are rare and thus it is dif cult to determine how subjects would respond to receiving a low offer.

To analyze the results, we use linear regression with robust SEs. We nd a signi cant negative effect of clarity of the successfulness of others on UG offer (coeff = −1.327, P = 0.016), as well as a signi cant negative effect of the consistency of others on UG offer (coeff = −1.241, P = 0.031). Similarly, we nd a signi cant negative effect of clarity of the successfulness of others on UG demand (coeff = −1.097, P = 0.029), but no signi cant effect of the consistency of others on UG demand (coeff = 0.294, P = 0.568). These results are qualitatively unchanged when including controls for age, sex, income, education and US residency (clarity of success predicting offer: coeff = −1.412, P = 0.024; consistency predicting offer: coeff = −1.143, P = 0.038; clarity of success predicting demand: coeff = −1.102, P = 0.044; consistency predicting demand: coeff = 0.100, P = 0.854). This experiment was approved by the Harvard University Committee on the Use of Human Subjects in Research, Application F17468.

1. Camerer CF (2003) Behavioral Game Theory: Experiments in Strategic Interaction (Princeton Univ Press, Princeton, NJ).
2. Güth W, Schmittberger R, Schwarze B (1982) An experimental analysis of ultimatum bargaining. J Econ Behav Organ 3(4):367–388.
3. Larrick RP, Blount S (1997) The claiming effect: Why players are more generous in social dilemmas than in ultimatum games. J Pers Soc Psychol 72(4):810–825.
4. Wallace B, Cesarini D, Lichtenstein P, Johannesson M (2007) Heritability of ultimatum game responder behavior. Proc Nat Acad Sci USA 104(40):15631–15634.
5. Zak PJ, Stanton AA, Ahmadi S (2007) Oxytocin increases generosity in humans. PLoS ONE 2(11):e1128.
6. Solnick S (2001) Gender differences in the ultimatum game. Econ Inq 39(2):189–200.
7. Kahneman D, Knetsch JL, Thaler RH (1986) Fairness and the assumptions of economics. J Business 59(4):S285–S300.
8. Straub PG, Murnighan JK (1995) An experimental investigation of ultimatum games: information, fairness, expectations, and lowest acceptable offers. J Econ Behav Organ 27(3):345–364.
9. Fehr E, Schmidt K (1999) A theory of fairness, competition and cooperation. Q J Econ 114(3):817–868.
10. Dawes CT, Fowler JH, Johnson T, McElreath R, Smirnov O (2007) Egalitarian motives in humans. Nature 446(7137):794–796.
11. Blake PR, Rand DG (2010) Currency value moderates equity preference among young children. Evol Hum Behav 31(3):210–218.
12. Fehr E, Bernhard H, Rockenbach B (2008) Egalitarianism in young children. Nature 454(7208):1079–1083.
13. Blake PR, McAuliffe K (2011) "I had so much it didn't seem fair": Eight-year-olds reject two forms of inequity. Cognition 120(2):215–224.
14. Shaw A, Olson KR (2012) Children discard a resource to avoid inequity. J Exp Psychol Gen 141(2):382–395.
15. Roth AE, Prasnikar V, Okuno-Fujiwara M, Zamir S (1991) Bargaining and market behavior in Jerusalem, Ljubljana, Pittsburgh, and Tokyo: An experimental study. Am Econ Rev 81(5):1068–1095.
16. Wells JS, Rand DG (2012) Strategic self-interest can explain seemingly "fair" offers in the Ultimatum Game. Available at http://ssrn.com/abstract=2136707. Accessed January 3, 2012.
17. Lin H, Sunder S (2002) Using experimental data to model bargaining behavior in Ultimatum Games. Experimental Business Research, ed Rapoport A (Kluwer, Dordrecht), pp 373–397.
18. Henrich J, et al. (2001) In search of Homo Economicus: Behavioral experiments in 15 small-scale societies. Am Econ Rev 91(2):73–78.
19. Hofbauer J, Sigmund K (1998) Evolutionary Games and Population Dynamics (Cambridge Univ Press, Cambridge, UK).
20. Nowak MA, Sigmund K (2004) Evolutionary dynamics of biological games. Science 303(5659):793–799.
21. Maynard Smith J (1982) Evolution and the Theory of Games (Cambridge Univ Press, Cambridge, UK).
22. Weibull J (1997) Evolutionary Game Theory (MIT Press, Cambridge, MA).
23. Henrich J, et al. (2005) "Economic man" in cross-cultural perspective: Behavioral experiments in 15 small-scale societies. Behav Brain Sci 28(6):795–815, discussion 815–855.
24. Traulsen A, Hauert C, De Silva H, Nowak MA, Sigmund K (2009) Exploration dynamics in evolutionary games. Proc Natl Acad Sci USA 106(3):709–712.
25. Young HP (1993) An evolutionary model of bargaining. J Econ Theory 59(1):145–168.
26. Manapat ML, Rand DG, Pawlowitsch C, Nowak MA (2012) Stochastic evolutionary dynamics resolve the Traveler's Dilemma. J Theor Biol 303:119–127.
27. Taylor P, Jonker L (1978) Evolutionary stable strategies and game dynamics. Math Biosci 40(1–2):145–156.
28. Hofbauer J, Schuster P, Sigmund K (1979) A note on evolutionary stable strategies and game dynamics. J Theor Biol 81(3):609–612.
29. Hofbauer J, Sigmund K (1990) Adaptive dynamics and evolutionary stability. Appl Math Lett 3(4):75–79.
30. Nowak MA, Page KM, Sigmund K (2000) Fairness versus reason in the ultimatum game. Science 289(5485):1773–1775.
31. Kandori M (1992) Social norms and community enforcement. Rev Econ Stud 59:63–80.
32. Nowak MA, Sigmund K (1998) Evolution of indirect reciprocity by image scoring. Nature 393(6685):573–577.
33. Nowak MA, Sigmund K (2005) Evolution of indirect reciprocity. Nature 437(7063): 1291–1298.
34. Ohtsuki H, Iwasa Y (2006) The leading eight: Social norms that can maintain cooperation by indirect reciprocity. J Theor Biol 239(4):435–444.
35. Pfeiffer T, Tran L, Krumme C, Rand DG (2012) The value of reputation. J R Soc Interface 9(76):2791–2797.
36. Manapat ML, Nowak MA, Rand DG (2012) Information, irrationality and the evolution of trust. J Econ Behav Organ, 10.1016/j.jebo.2012.10.018.
37. Güth W, Yaari M (1992) An evolutionary approach to explain reciprocal behavior in a simple strategic game. Explaining Process and Change - Approaches to Evolutionary Economics, ed Witt U (Univ of Michigan Press, Ann Arbor), pp 23–34.
38. Ellingsen T (1997) The evolution of bargaining behavior. Q J Econ 112(2):581–602.
39. Gale J, Binmore KG, Samuelson L (1995) Learning to be imperfect: The ultimatum game. Games Econ Behav 8(1):56–90.
40. Huck S, Oechssler J (1999) The indirect evolutionary approach to explaining fair allocations. Games Econ Behav 28:13–24.
41. Boyd R, Richerson PJ (1990) Group selection among alternative evolutionarily stable strategies. J Theor Biol 145(3):331–342.
42. Traulsen A, Nowak MA (2006) Evolution of cooperation by multilevel selection. Proc Natl Acad Sci USA 103(29):10952–10955.
43. Nowak MA, Sasaki A, Taylor C, Fudenberg D (2004) Emergence of cooperation and evolutionary stability in nite populations. Nature 428(6983):646–650.
44. Simon HA (1972) Theories of bounded rationality. Decisions and Organisation, eds McGuire CB, Radner R (North-Holland, Amsterdam).
45. Rand DG, Nowak MA (2012) Evolutionary dynamics in nite populations can explain the full range of cooperative behaviors observed in the centipede game. J Theor Biol 300:212–221.
46. Ewens WJ (2004) Mathematical Population Genetics. I. Theoretical Introduction (Springer, New York).
47. Moran PAP (1962) The Statistical Processes of Evolutionary Theory (Clarendon, Oxford).
48. Fudenberg D, Imhof LA (2006) Imitation processes with small mutations. J Econ Theory 131(1):251–262.
49. Hauert C, Traulsen A, Brandt H, Nowak MA, Sigmund K (2007) Via freedom to coercion: The emergence of costly punishment. Science 316(5833):1905–1907.
50. Hauert C, Traulsen A, Brandt H, Nowak MA, Sigmund K (2008) Public goods with punishment and abstaining in nite and in nite populations. Biol Theory 3(2): 114–122.
51. De Silva H, Hauert C, Traulsen A, Sigmund K (2010) Freedom, enforcement, and the social dilemma of strong altruism. J Evol Econ 20(2):203–217.
52. Antal T, Traulsen A, Ohtsuki H, Tarnita CE, Nowak MA (2009) Mutation-selection equilibrium in games with multiple strategies. J Theor Biol 258(4):614–622.
53. Page KM, Nowak MA (2001) A generalized adaptive dynamics framework can describe the evolutionary Ultimatum Game. J Theor Biol 209(2):173–179.
54. Horton JJ, Rand DG, Zeckhauser RJ (2011) The online laboratory: Conducting experiments in a real labor market. Exp Econ 14(3):399–425.
55. Rand DG (2012) The promise of Mechanical Turk: How online labor markets can help theorists run behavioral experiments. J Theor Biol 299:172–179.
56. Mason W, Suri S (2012) Conducting behavioral research on Amazon's Mechanical Turk. Behav Res Methods 44(1):1–23.
57. Amir O, Rand DG, Gal YK (2012) Economic games on the internet: The effect of $1 stakes. PLoS ONE 7(2):e31461.
58. Paolacci G, Chandler J, Ipeirotis PG (2010) Running experiments on Amazon Mechanical Turk. Judgm Decis Mak 5(5):411–419.
59. Buhrmester MD, Kwang T, Gosling SD (2011) Amazon's Mechanical Turk: A new source of inexpensive, yet high-quality, data? Perspect Psychol Sci 6(1):3–5.
60. Rand DG, Arbesman S, Christakis NA (2011) Dynamic social networks promote cooperation in experiments with humans. Proc Natl Acad Sci USA 108(48):19193–19198.
61. Rand DG, Nowak MA (2011) The evolution of antisocial punishment in optional public goods games. Nat Commun 2:434.
62. Pillutla MM, Murnighan JK (1996) Unfairness, anger, and spite: Emotional rejections of ultimatum offers. Organ Behav Hum Decis Process 68(3):208–224.
63. Sanfey AG, Rilling JK, Aronson JA, Nystrom LE, Cohen JD (2003) The neural basis of economic decision-making in the Ultimatum Game. Science 300(5626):1755–1758.
64. Knoch D, Pascual-Leone A, Meyer K, Treyer V, Fehr E (2006) Diminishing reciprocal fairness by disrupting the right prefrontal cortex. Science 314(5800):829–832.
65. Tabibnia G, Satpute AB, Lieberman MD (2008) The sunny side of fairness: Preference for fairness activates reward circuitry (and disregarding unfairness activates self-control circuitry). Psychol Sci 19(4):339–347.
66. Rand DG, Greene JD, Nowak MA (2012) Spontaneous giving and calculated greed. Nature 489(7416):427–430.

Supporting Information

Rand et al. 10.1073/pnas.1214167110

SI Text

We study the evolutionary dynamics of the Ultimatum Game (UG) analytically in two ways. The first approach, described in SI Single-Population Formation, is a "single-population" formulation: we consider a population of size N in which each individual can be both a proposer and a responder with equal probability. The second approach, described in SI Two-Population Formation, is a "two-population" formulation: we consider two populations of size N, one of which is a population of proposers and the other, a population of responders, and we study their coevolution. Both approaches will yield the same result; an intuitive explanation of this result is then provided in SI Intuition on the Role of Mutation.

We then turn to additional agent based simulation results. Fig. S1 shows that the results in Fig. 1 in the main text are robust to different mutation rates. SI Local Mutation Kernel and Figs. S2 and S3 describe our simulations using a local mutation kernel (rather than the global mutation kernel used in the main text), and show that the results in Fig. 2 in the main text are robust to this alternate mutation structure. Fig. S4 shows the average values of p and q from numerous behavioral experiments as well as from agent-based simulations using particular parameter sets highlighted in Fig. 2 in the main text. Fig. S5 shows that the results in Fig. 2 in the main text are robust to different population sizes. Fig. S6 shows the distribution of individual-level offers and demands for a representative set of parameters.

SI Single-Population Formulation

Here, we consider a population of N players who play the role of proposers and responders with equal probability. We specify an agent's strategy as the pair $S = (p; q)$, $0 \leq p \leq 1$, and $0 \leq q \leq 1$, where p is the amount offered when acting as proposer, and q is the minimum amount demanded when acting as responder, or the rejection threshold. Hence, the strategy space for the UG is the unit square. Let $A(S_1; S_2)$ be the expected payoff that strategy $S_1 = (p_1; q_1)$ gets from strategy $S_2 = (p_2; q_2)$. Because we assume that in the interaction between a player using strategy S_1 and a player using strategy S_2 each player can be in the role of the proposer with equal probability, $A(S_1; S_2)$ is given (up to a $1/2$ factor which we henceforth omit) by the function

$$A(S_1; S_2) = \begin{cases} 1 - p_1 + p_2 & \text{if } p_1 \geq q_2 \text{ and } p_2 \geq q_1 \\ 1 - p_1 & \text{if } p_1 \geq q_2 \text{ and } p_2 < q_1 \\ p_2 & \text{if } p_1 < q_2 \text{ and } p_2 \geq q_1 \\ 0 & \text{if } p_1 < q_2 \text{ and } p_2 < q_1 \end{cases} \quad [S1]$$

Every individual in the population plays the UG with every other individual and they all get payoffs according to the function above. The (relative) fecundity (or effective payoff) of a player with an average payoff π is given by $\exp(w\pi)$, where the parameter $w > 0$ represents the intensity of selection. Individuals reproduce proportional to their fecundity. In each time step a random player dies and another player (including the dying player himself) is picked proportional to fecundity to replace the dead. Reproduction is subject to mutation: the offspring inherits the strategy of the parent with probability

1 – u and with probability u it adopts a strategy uniformly at random. We say that strategy S is favored on average in the mutation-selection equilibrium, if its abundance exceeds the mean.

Let us first assume that our strategies do not cover the entire unit square, but in fact are only of the form $s = (i=m; j=m)$ with $1 \leq i; j \leq m$ being integers. This discretizes the problem, making it possible to invoke previous results. To then go back to the continuous strategy space we simply take the limit $m \to \infty$.

Having turned our continuous problem into a discrete one, we are now interested in the stationary abundance of these discrete strategies. For this problem, we can use the result in ref. 1 to conclude that, for large population size N, strategy s is favored by selection if $L_s + NuH_s > 0$, where

$$L_s = \frac{1}{m^2} \sum_{i=1}^{m} \sum_{j=1}^{m} f A(s; s') + A(s'; s) - A(s'; s') - A(s'; s') g$$

$$H_s = \frac{1}{m^4} \sum_{i=1}^{m} \sum_{j=1}^{m} \sum_{i'=1}^{m} \sum_{j'=1}^{m} f A(s; s') - A(s'; s') g \quad [S2]$$

Here, $s' = (i'=m; j'=m)$ and $s'' = (i''=m; j''=m)$. Moreover, ref. 1 showed that the higher the quantity $L_s + NuH_s > 0$, the more the strategy s is favored by selection. Consequently, to determine which strategy is most favored by selection, one simply has to maximize $L_s + NuH_s > 0$.

Taking the limit $m \to \infty$ as in Tarnita et al. (2), the sums in S2 converge to the integrals

$$\mathcal{L}_S = \int_0^1 \int_0^1 A(S; S') + A(S'; S) - A(S'; S') - A(S'; S') \, dp' \, dq'$$

$$\mathcal{H}_S = \int_0^1 \int_0^1 \int_0^1 \int_0^1 A(S'; S'') - A(S'; S'') \, dp' \, dq' \, dp'' \, dq'' ; \quad [S3]$$

where $S' = (p'; q')$ and $S'' = (p''; q'')$. Moreover, it follows that the condition for strategy S to be favored by selection is $\mathcal{L}_S + Nu\mathcal{H}_S > 0$ and that the most favored strategy is determined by maximizing $\mathcal{L}_S + Nu\mathcal{H}_S > 0$. Depending on whether $p \geq q$ or $p < q$, the payoff function $A(S; S')$ takes two different values and hence we find the condition for strategy S to be favored by selection to be

$$\mathcal{L}_S + Nu\mathcal{H}_S = I(p \geq q) + p^2(-Nu - 2) + q^2 \left(-\frac{Nu}{2} - 1\right)$$

$$+ p(1 + Nu) + q - \frac{1}{2} > 0; \quad [S4]$$

where $I(\text{condition})$ is 1 if condition is true and is 0 if condition is false.

Maximizing S4, we conclude that the optimum strategy (most abundant in the stationary distribution and hence, by our measure, most favored by selection) is achieved when $p \geq q$ and is given by

$$\left(p_{opt}, q_{opt}\right) = \begin{cases} \left(\dfrac{1}{3}, \dfrac{1}{3}\right) & \text{if } 0 \leq Nu \leq 1 \\[2mm] \left(\dfrac{1+Nu}{4+2Nu}, \dfrac{1}{2+Nu}\right) & \text{if } Nu > 1 \end{cases}.$$

Note that for low mutation, the optimum strategy is $(1/3, 1/3)$. Hence, the most successful strategy is one that offers 33% and also rejects any offer lower than 33%. As mutation increases, the proposal increases and the rejection threshold decreases. For high mutation, the most frequent strategy is $(1/2, 0)$; thus the proposal is 50% and the rejection threshold is 0.

SI Two-Population Formulation

Next we will derive the same results as above but using a different approach. Instead of considering a population where each individual can be both proposer and responder, we consider two distinct populations—the population of proposers and that of responders—and explore their evolutionary game dynamics. This means that, unlike in SI Single-Population Mutation, where the strategy of an individual was given by a vector $S = (p, q) \in [0,1] \times [0,1]$, here the strategy of an individual is given by one number. Thus, an individual from the population of proposers will have strategy $S_{prop} = p \in [0,1]$, which represents the offer he makes and an individual from the population of responders will have strategy $S_{resp} = q \in [0,1]$, which represents his rejection threshold. When two such players meet, their payoffs are given by

$$A_{prop}(p, q) = \begin{cases} 1 - p & \text{if } p \geq q \\ 0 & \text{if } p < q \end{cases},$$
$$A_{resp}(p, q) = \begin{cases} p & \text{if } p \geq q \\ 0 & \text{if } p < q \end{cases}; \qquad [S5]$$

where A_{prop} and A_{resp}, respectively, represent payoffs of the proposer and the responder.

Suppose that there are N players in each population (i.e., N proposers and N responders). Each proposer plays the UG described above with every responder in the responders' population and obtains an average game payoff. Similarly, each responder plays the game with every proposer and obtains an average game payoff. We assume that selection occurs in each population according to payoffs in the UG. More specifically, a random player in either population dies (in cultural evolution terms, he attempts to change his strategy) and another player in the same population (it can be the dying player himself) replaces the dead with its offspring with probability proportional to one's fecundity. We assume that one's (relative) fecundity is given by $\exp[w\pi]$, where π represents one's average payoff in the game and $w > 0$ represents the intensity of selection. Reproduction (imitation) occurs with mistakes—with probability u, the offspring is susceptible to mutation and randomly adopts a strategy uniformly at random, independently of its parent's strategy.

For simplicity of our analysis, we will first discretize the problem, as before. Thus, we first consider that the possible proposals have the form $s_{prop} = i/m$ and the possible rejection thresholds have the form $s_{resp} = j/m$, where $m \geq 1$ is an integer and $1 \leq i, j \leq m$. In this case, assuming weak selection, $w \to 0$, Ohtsuki et al. (3) have obtained the result that the combination of proposer's and responder's strategies that is most abundant in the stationary distribution (and hence, by our definition, is most favored by selection) is the one that maximizes $L\left(\frac{i}{m}, \frac{j}{m}\right) + 2(N-1)u \cdot H\left(\frac{i}{m}, \frac{j}{m}\right)$, where

$$L\left(\frac{i}{m}, \frac{j}{m}\right) = \frac{1}{m^2}\sum_{i,j=1}^{M}\left[A_{prop}\left(\frac{i}{m}, \frac{j}{m}\right) - A_{prop}\left(\frac{i}{m}, \frac{j}{m}\right)\right]$$
$$+ A_{prop}\left(\frac{i}{m}, \frac{j}{m}\right) - A_{prop}\left(\frac{i}{m}, \frac{j}{m}\right)$$
$$+ \frac{1}{m^2}\sum_{i,j=1}^{M}\left[A_{resp}\left(\frac{i}{m}, \frac{j}{m}\right) - A_{resp}\left(\frac{i}{m}, \frac{j}{m}\right)\right] \qquad [S6]$$
$$+ A_{resp}\left(\frac{i}{m}, \frac{j}{m}\right) - A_{resp}\left(\frac{i}{m}, \frac{j}{m}\right)$$

$$H\left(\frac{i}{m}, \frac{j}{m}\right) = \frac{1}{m^2}\sum_{i,j=1}^{M}\left[A_{prop}\left(\frac{i}{m}, \frac{j}{m}\right) - A_{prop}\left(\frac{i}{m}, \frac{j}{m}\right)\right]$$
$$+ A_{resp}\left(\frac{i}{m}, \frac{j}{m}\right) - A_{resp}\left(\frac{i}{m}, \frac{j}{m}\right)$$

A direct calculation shows

$$L\left(\frac{i}{m}, \frac{j}{m}\right) = I(i \geq j) - 2\left(\frac{i}{m}\right)^2 - \left(\frac{j}{m}\right)^2 + \frac{i}{m} + \frac{j}{m} - \frac{1}{2} + \frac{j}{m^2} - \frac{1}{2m}$$

$$H\left(\frac{i}{m}, \frac{j}{m}\right) = -\left(\frac{i}{m}\right)^2 - \frac{1}{2}\left(\frac{j}{m}\right)^2 + \frac{i}{m} + \frac{j}{2m^2}; \qquad [S7]$$

where $I(\text{condition})$ is 1 if condition is true and is 0 if condition is false. Now let $p = i/m$ and $q = j/m$. Substituting $i = pm$ and $j = qm$ and taking the limit $m \to \infty$ gives

$$L(p, q) = I(p \geq q) - 2p^2 - q^2 + p + q - \frac{1}{2}$$
$$H(p, q) = -p^2 - \frac{1}{2}q^2 + p \qquad [S8]$$

A direct calculation now shows that, for large N, the most abundant pair of strategies (p_{opt}, q_{opt}) [obtained by maximizing $L(p, q) + 2(N-1)uH(p, q)$] is given by

$$\left(p_{opt}, q_{opt}\right) = \begin{cases} \left(\dfrac{1}{3}, \dfrac{1}{3}\right) & \text{if } 0 < 2Nu \leq 1 \\[2mm] \left(\dfrac{1+2Nu}{4+4Nu}, \dfrac{1}{2+2Nu}\right) & \text{if } 1 < 2Nu \end{cases} \qquad [S9]$$

This result is exactly the same as the one for the single-population model in SI Single-Population Mutation except that Nu is replaced with 2Nu. This makes sense, because the total population size in the two-population formulation is $N + N = 2N$, whereas it is N in the one-population formulation.

SI Intuition on the Role of Mutation

As we have seen, the magnitude of the rescaled mutation rate Nu is a crucial determinant of the strategy that is favored the most by natural selection (= the one that is most frequently observed in the mutation-selection equilibrium of our stochastic evolutionary dynamics). Here, we try to explain the reason for that.

In our formulation, a mutant strategy, $S = (p, q)$, is randomly chosen from our strategy space, which is the unit square $[0,1] \times [0,1]$. This assumption means that a mutant almost surely adopts a strategy that is not observed in a current population. Therefore, mutation increases variation in strategies. At the same time, finiteness of the population size reduces the variation via random sampling of a finite number of offspring. What is then important is how many different strategies coexist in a pop-

ulation of size N. Because our model corresponds to Moran infinitely many alleles model in population genetics, many results are already known. The expected number of different strategies coexisting in the population is given by the following exact formula (4):

$$\sum_{k=0}^{N-1} \frac{Nu}{Nu + k(1-u)}. \qquad [\text{S10}]$$

Table S1 shows some values of this expression for various N and u. From S10 (but also from Table S1) we see that, in the low mutation limit $Nu \to 0$, the number of different strategies in the population is close to 1, suggesting that the population is almost always monomorphic. When a new and rare mutant appears in the population, the number of different strategies becomes 2 (resident and mutant), and we expect that mutants will either die out or take over the resident population before another new mutant arises. Thus, at most two strategies are involved in a takeover at any moment in time. Therefore, in the low mutation, a strategy is selected if it can resist invasion by a single randomly chosen strategy. Hence, what needs to be maximized is expected relative payoff in pairwise competition with a random opponent.

As Nu increases, the number of different strategies present in the population also increases. In the high mutation limit, all strategies are present in the population simultaneously with approximately equal frequency. Hence, the optimum strategy is the one that maximizes its expected absolute payoff against a randomly chosen opposing strategy. At intermediate mutation rates, the evolutionary dynamic has characteristics of both the fully heterogeneous and fully homogeneous extremes. How these conclusions play a role in determining the winning strategies is explained at length in the main text.

SI Local Mutation Kernel

In the main text analysis and simulations, a mutant's p and q values are randomly picked from the uniform distribution $[0;1]$. Thus, mutations are "global," in the sense that a mutant's new strategy has no relation to the previous strategy. An alternative scheme, however, uses local mutation, where the mutant strategy is some perturbation off of the parent strategy. To investigate the effects of local mutation, we use the following mutation kernel. For a parent with $p = p^*$, the mutant p is picked from a β-distribution defined by the probability density function

$$f(p) = \frac{p^{\alpha-1}(1-p)^{\beta-1}}{\int_0^1 u^{\alpha-1}(1-u)^{\beta-1}du}; \qquad [\text{S11}]$$

where $\alpha = (p^{*2} - 2p^* + 1)/(1-p^*)$ and $\beta = \kappa$ if $x < 0.5$ or $\beta = \kappa$ and $\alpha = ((1-p^*)^2 - 2(1-p^*) + 1)/p^*$ if $x \geq 0.5$, and κ is a parameter determining how similar the mutant tends to be to the parent. A β-distribution is used as this distribution is bounded on the interval $[0;1]$ and is unimodal if $\alpha; \beta > 1$. The particular values of α and β are chosen such that the modal value of the distribution is p^*.

Fig. S2 shows sample probability density functions for different values of p^*, using $\kappa = 50$. The same distribution is used to independently generate q values. Fig. S3 shows the results of repeating the simulations shown in Fig. 2 in the main text, but now using this local mutation kernel $\kappa = 50$. We see qualitative agreement: Across a wide range of w and u values, we observe average q > 0 and p > q. Thus, our results are robust to the use of a local mutation kernel.

1. Antal T, Traulsen A, Ohtsuki H, Tarnita CE, Nowak MA (2009) Mutation-selection equilibrium in games with multiple strategies. J Theor Biol 258(4):614–622.
2. Tarnita CE, Antal T, Nowak MA (2009) Mutation-selection equilibrium in games with mixed strategies. J Theor Biol 261(1):50–57.
3. Ohtsuki H (2010) Stochastic evolutionary dynamics of bimatrix games. J Theor Biol 264(1):136–142.
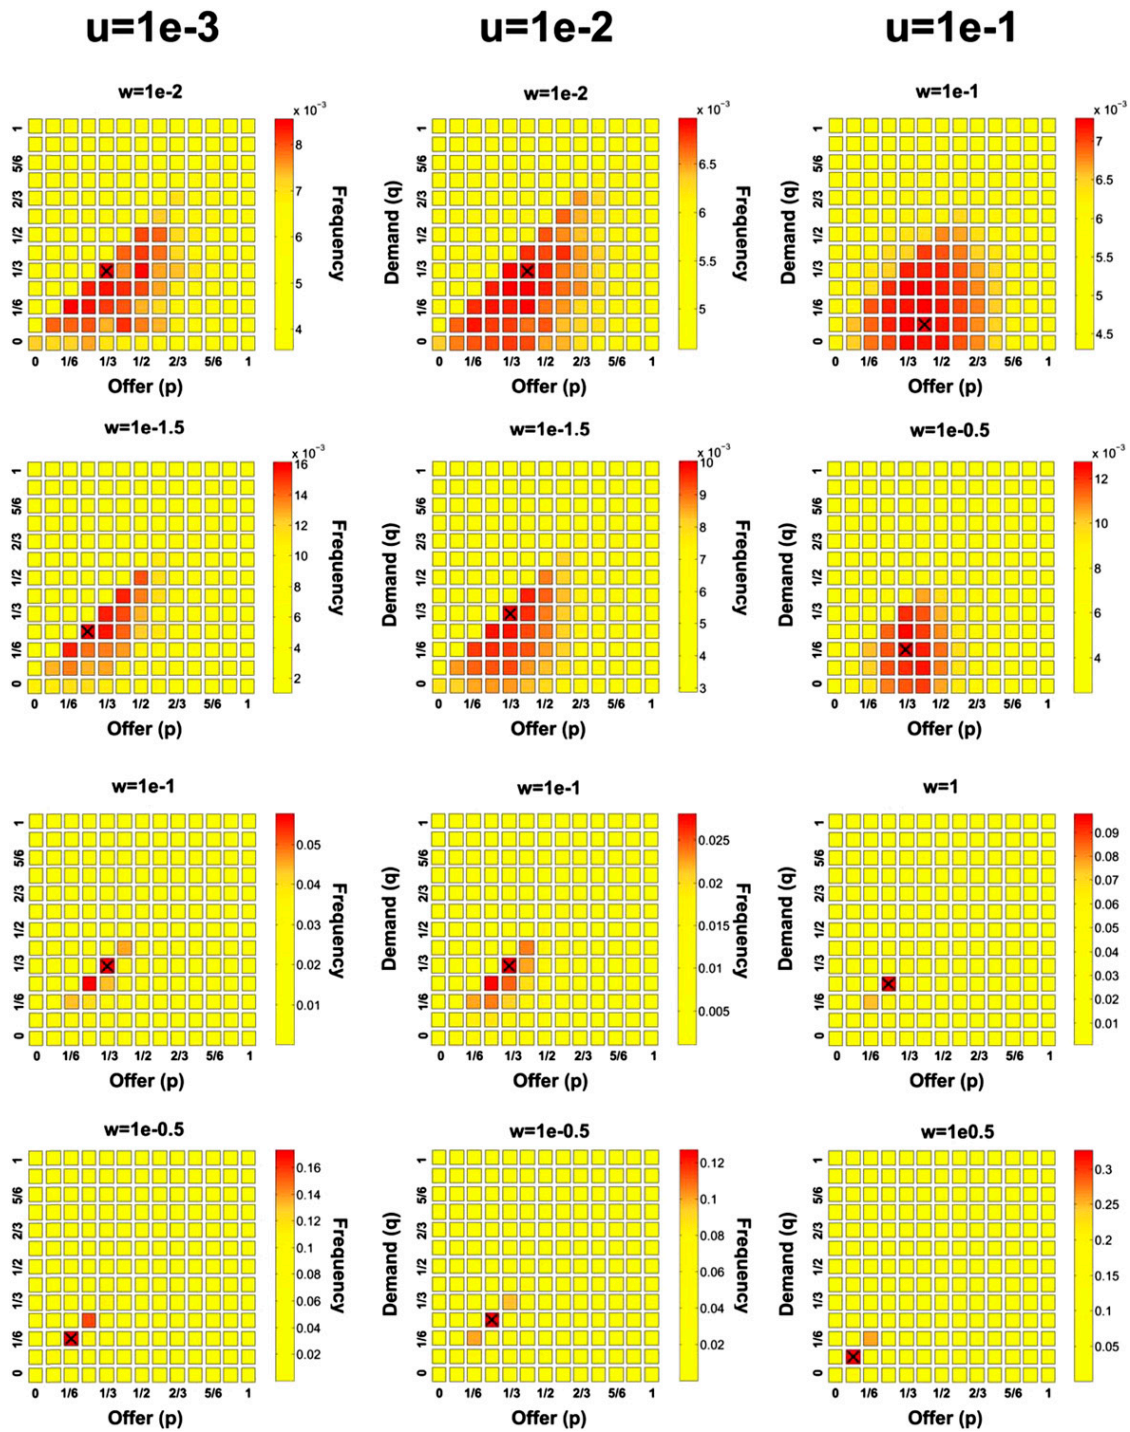4. Ewens WJ (2004) Mathematical Population Genetics: I. Theoretical Introduction (Springer, New York).

Fig. S1. For u = 10⁻³, u = 10⁻², and u = 10⁻¹, respectively, and for N = 100, we show frequencies of {p; q} pairs over 10⁸ generations, binned in increments of 0.1. The most common bin is indicated with a black ×.
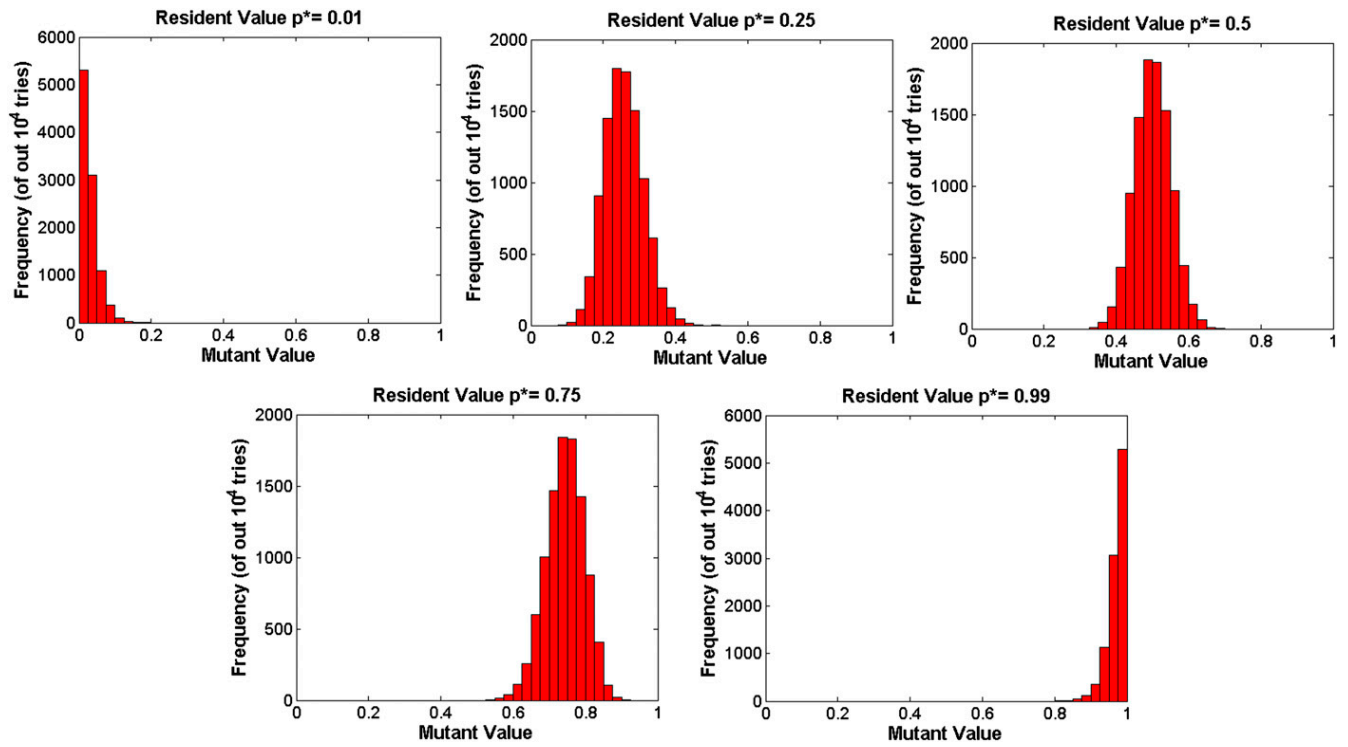
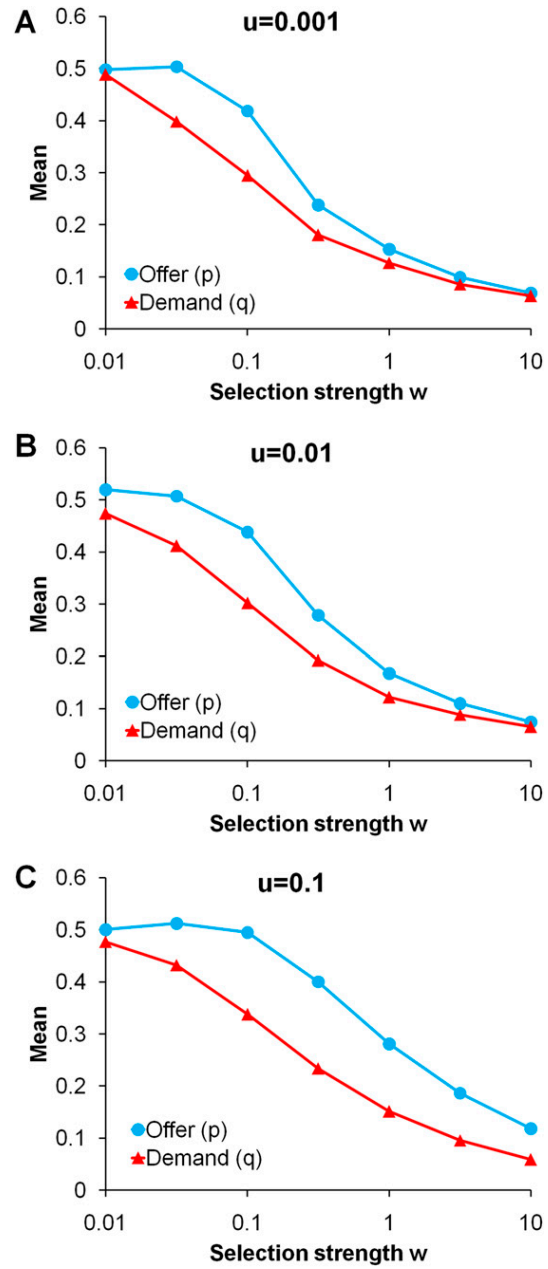Fig. S2. Local mutation kernel probability density functions for parent p*, using = 50.

Fig. S3.   Agent-based simulation results using local mutation with = 50. Shown are time-averaged values of p and q over 10^8 generations, using population N = 100 and (A) u = 10^{-3}, (B) u = 10^{-2}, and (C) u = 10^{-1}.
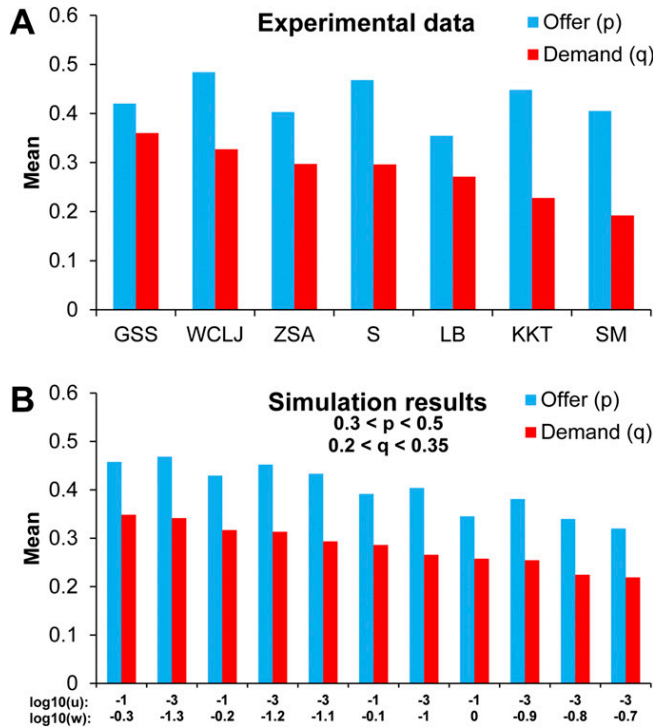
Fig. S4. Mean values of offers by player 1, p (blue), and minimum amount demanded by player 2, q (red), are shown from data (A) and evolutionary agent-based simulations (B) Observations are sorted by q value. (A) Data as presented in refs. 1–7. (B) All simulations with 0:3 < p < 0:5 and 0:2 < q < 0:35 are shown, from the set of mutation rate u = ⅟10⁻³; 10⁻¹ and selection strength w in log-scaled increments of 0.1. Simulations use population size N = 100, and show time-averaged values of p and q over 10^8 generations. Fig. 2 in main text shows a systemic exploration of the relationship between u, w, and the average strategy.

1. Güth W, Schmittberger R, Schwarze B (1982) An experimental analysis of ultimatum bargaining. J Econ Behav Organ 3(4):367–388.
2. Wallace B, Cesarini D, Lichtenstein P, Johannesson M (2007) Heritability of ultimatum game responder behavior. Proc Nat Acad Sci USA 104(40):15631–15634.
3. Zak PJ, Stanton AA, Ahmadi S (2007) Oxytocin increases generosity in humans. PLoS ONE 2(11):e1128.
4. Solnick S (2001) Gender differences in the ultimatum game. Econ Inq 39(2):189–200.
5. Larrick RP, Blount S (1997) The claiming effect: Why players are more generous in social dilemmas than in ultimatum games. J Pers Soc Psychol 72(4):810–825.
6. Kahneman D, Knetsch JL, Thaler RH (1986) Fairness and the assumptions of economics. J Business 59(4):S285–S300.
7. Straub PG, Murnighan JK (1995) An experimental investigation of ultimatum games: information, fairness, expectations, and lowest acceptable offers. J Econ Behav Organ 27(3): 345–364.
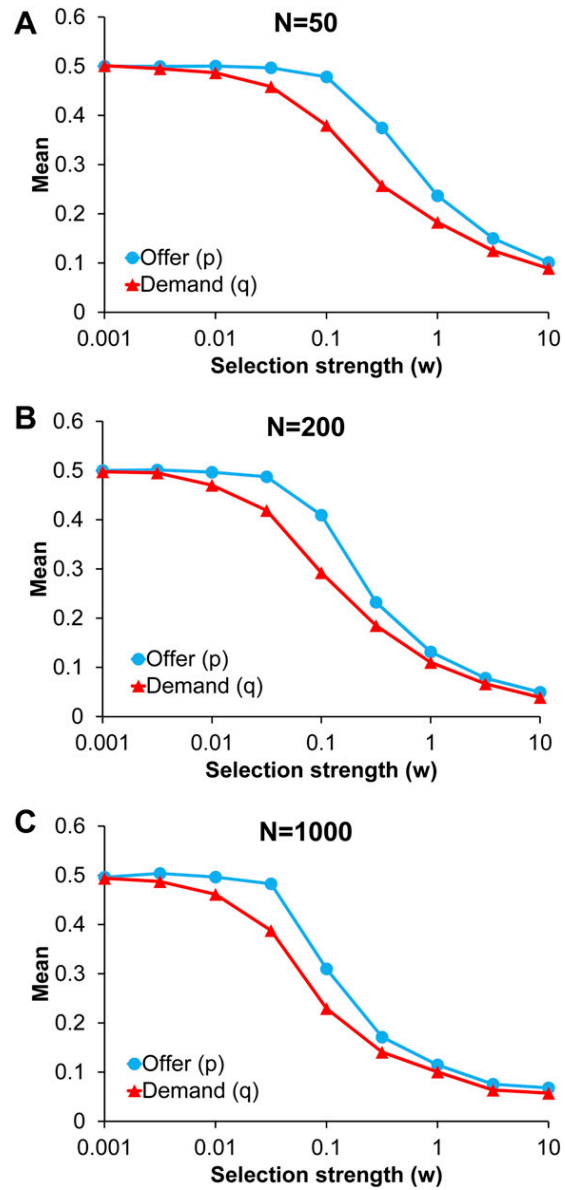
Fig. S5. Agent-based simulation results using global mutation (as in the main text) but varying the population size N. Shown are time-averaged values of p and q over 10^8 generations, using mutation rate u = 0:01 and (A) N = 50, (B) N = 200, and (C) N = 1000.
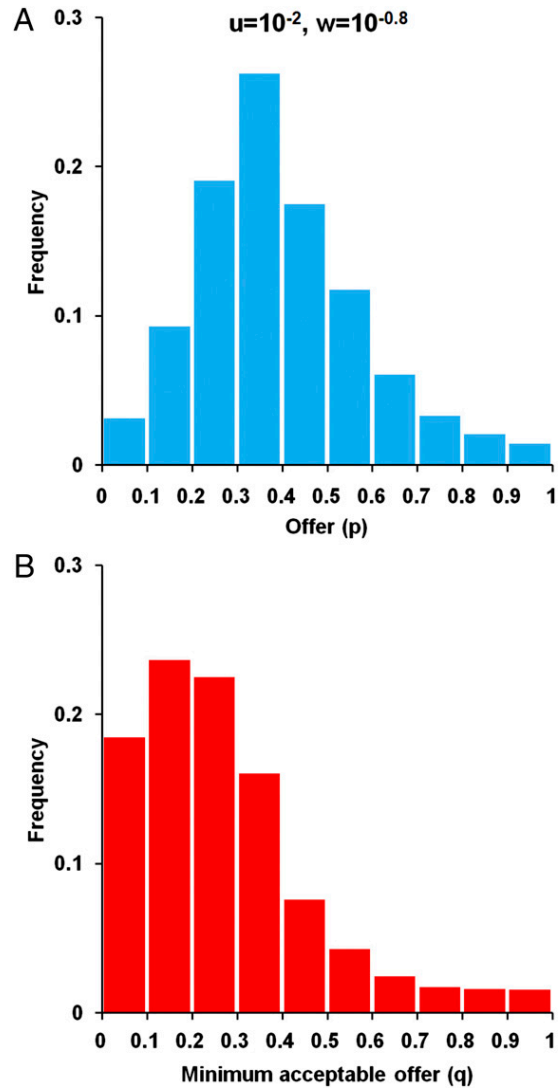
A $u=10^{-2}$, $w=10^{-0.8}$

B

Fig. S6.   Distribution of offers p (A) and demands q (B) for N = 100, u = 0:01, and w = $10^{-0:8}$. This is an alternative method of presenting the data shown in Fig. 1 in main text: here, p and q distributions are shown separately, rather than showing the joint ♮p; q distribution as is done in Fig. 1 in main text.

Table S1.   Expected number of different strategies coexisting in the population

Nu	N = 100	N = 1; 000	N = 10; 000
10	26.0417	47.0201	69.6545
1	5.22322	7.49132	9.78842
0.1	1.50298	1.73318	1.96342
0.01	1.05162	1.07468	1.09771

A random world is a fair world

James H. Fowler[a,b,1] and Nicholas A. Christakis[c,d]
[a]Medical Genetics Division, School of Medicine, University of California at San Diego, La Jolla, CA 92093; [b]Political Science Department, University of California at San Diego, La Jolla, CA 92093; [c]Department of Sociology, Harvard University, Cambridge, MA 02138; and [d]Department of Medicine, Harvard Medical School, Boston, MA 02115

A preference for fairness or equity in the distribution of resources influences many human decisions (1). The origin of this preference is a topic that has consumed philosophers (2), social scientists (3), and biologists (4) for centuries. However, although we feel a sense of fairness deeply and intuitively, it has so far been difficult to explain from first principles how such a feeling might have evolved. How could natural selection allow for the survival of "fair" individuals who sometimes give things away to equalize resources when they must compete with self-interested individuals who keep everything for themselves? In PNAS, Rand et al. (5) provide a unique and compelling solution to this puzzle: it's all because of dumb luck.

To study fairness, authors use the so-called "ultimatum game" (6). In this game, one person (the proposer) offers a specific division of a sum of money, and the other (the responder) decides whether to accept this offer. If the responder accepts, they each receive the amount of money as proposed. If the responder rejects the offer, they each receive nothing. If both players are rational and self-interested and they play the game only once, then the responder should accept any nonzero offer (something is better than nothing!). Knowing that, the proposer should offer slightly more than zero to the responder and keep the rest for himself.

However, this result is not what we observe, anywhere. Dozens of studies in both large (7) and small-scale (8) societies show that proposers tend to make "fair" offers, in the range of 30–50% for the responder. Furthermore, responders tend to demand such behavior, rejecting offers when they fall below 20–35%.

Past efforts to explain the origin of these preferences have used deterministic game theory, which assumes that individuals with higher expected payoffs will always come to dominate the population (9). These models cannot explain fair offers or rejection of nonzero offers without making additional assumptions. For example, if we assume individuals have information about others' past behavior, then they can make strategies contingent on the reputation of their opponents, and this will benefit individuals with a reputation for rejecting low offers. However, how did the individuals get this information? And how do they avoid being exploited by individuals who can fake such a reputation? The additional assumptions are complex and hard to justify.

Instead, Rand et al. (5) return to first principles and use a different approach. Rather than assuming that evolution is deterministic, they assume it is stochastic. In reality, evolution sometimes favors the lucky, especially when the relationship between payoffs and survival is weak. This theory means a variety of strategies can endure and the winning strategy must do well in such an environment. Intuitively, if some of the responders are rejecting nonzero offers—not because it is the best strategy but because it happens to survive sometimes—then proposers need to make fairer offers.

Proximate selfish behavior can be bad for you, and under evolutionary pressure may not even survive.

Remarkably, when Rand et al. (5) apply stochastic evolutionary game theory in this way, they find that offers exceed demands and demands are greater than zero, just as they are in the empirical data. This result is true under a very wide range of possible scenarios when they vary selection (the relationship between payoffs and survival) and mutation (the likelihood that an individual chooses a random strategy). In fact, in some of these scenarios, they can exactly reproduce the average offer and the average demand from experiments in behavioral economics.

If the article ended there it would already be impressive for the way this work explains the observed data with the most parsimonious model to date. However, Rand et al. (5) also use the model to make two unique predictions, both of which are confirmed by careful measurement in a sample of 140 subjects.

First, as selection becomes weaker, it increases the likelihood of survival for both proposers and responders who try alternative strategies. Therefore, people living in circumstances where it is harder to assess the success of others' strategies should make both higher offers and higher demands. Second, as the rate of mutation increases, it directly increases the average offer because the average without mutation is less than one-half. However, the effect of mutation on demands is more ambiguous. Random demands will tend to increase the average because they are also below one-half, but there is more to lose from rejection because the offers tend to be higher, and this favors lower demands. Therefore, people living in circumstances where others are inconsistent in their behavior should make higher offers but not necessarily have higher demands.

Why Randomness Matters

It may seem remarkable that randomness is what drives "fair" behavior in this model, but it is consistent with what we know about other human behaviors that apparently defy rational explanation: uncertainty is key.

For example, it is well known that humans tend to exhibit overconfidence. When making interpersonal comparisons on a wide variety of traits, most people think they are above average. Such a bias might cause individuals to engage in contests they are sure to lose. However, uncertainty about capabilities means that the overconfident are also more likely to win other contests because less-confident individuals may decide not to enter the contest in the first place. As a consequence, evolutionary models show that, counterintuitively, overconfidence maximizes individual fitness and populations tend to become overconfident under a wide variety of conditions (10). Similar behavior in the face of uncertainty is observed in physician prognostication (11).

Stochastic evolutionary game theory has not yet been widely used, but it is already yielding a variety of promising results, tting empirical data on human behavior better than deterministic models (12, 13). An important conceptual innovation in the stochastic theory is that selection is not the only important factor in evolution. Mutation is also important, and it is the balance between mutation and selection that ultimately matters for determining evolutionary outcomes. Although this may seem to complicate deterministic models, many basic models yield elegant closed-form solutions (14) and the analogs to agent-based evolutionary models are direct and easy to interpret.

For example, Nowak, Tarnita, and Wilson (15) recently applied stochastic evolutionary game theory to the problem of eusociality and showed that it could help to explain group-level evolutionary outcomes without any extra assumptions about "inclusive tness." Although dozens of other scholars wrote rebuttals (16) to this work (primarily to defend their use of more approximate models), these responses did not counter an important point: inclusive tness theory is a special case of a more general model that is simply based on individual selection under mutation and a precise elaboration of the set of interactions among individuals in the population. Thus, a random world is also one in which we can better understand how individual selection can drive group behavior.

Next Steps

An important challenge for stochastic game theory is whether or not it can be used to predict individual-level behavior. The model elaborated by Rand et al. (5) does an excellent job in matching population averages but, as they show in their supplementary information, there is wider variance in individual strategies than is normally present in empirical data. For example, their model yields too many individuals who make and demand offers above 50%.

Stochastic learning theory has faced similar challenges. Simple rules based on reinforcement learning (17) were used successfully to explain aggregate behavior in pigeons, gold-sh, and, in some situations, in humans, but they were abandoned by psychologists in the 1970s in part for their inability to predict individual-level behavior (7). However, this disconnect between group and individual results may be easy to x simply by adjusting a functional form in the model. For example, a simple model of voter behavior generates more realistic individual-level results when reinforcement yields xed percentage changes in behavior rather than changes that become smaller near-extreme values (18). Similarly, in the Rand et al. (5) model, it may be the case that local mutation yields less variance in individual behavior than global mutation, and in fact this may be a way to test what kinds of exploration strategies are most likely.

Given that the Rand et al. (5) model can be interpreted as either a learning or natural-selection model, it suggests a wide variety of possible mechanisms. Some of these mechanisms could be cultural, such as those advanced by Henrich et al. in their study of ultimatum game play in several small-scale societies (8). It would be interesting to conduct the same experiment used by Rand et al. in each of these societies to see if variation in expectations about successful opponents and the mutability of their game play could explain variation in mean offers and demands. Other work suggests that variation in ultimatum game play is heritable; in other words, genetic variation is, in part, driving the different strategies that people use when they play the ultimatum game (19). In addition, functional MRI studies of other behavioral games suggests that the ventromedial prefrontal cortex and the insula may play a mediating role between genes and a sense of fairness (20). Although the insula result has been interpreted in the context of its association with social decision-making, it would be interesting to see if the ventromedial prefrontal cortex activation is driven by processing expectations about others strategies.

Finally, although Rand et al. (5) generate their results with a model that is based on individual natural selection, it is fascinating that it yields behavior that may otherwise appear to be based on something else. Rand et al. call the process "self-interested natural selection" but later note that "myopic self-interest is vanquished whereas fairness triumphs." This is a nice turn of concept. Proximate sel sh behavior can be bad for you, and under evolutionary pressure may not even survive because fairness maximizes individual tness. It may not be fair to be sel sh, but it is certainly sel sh to be fair.

1 Dawes CT, Fowler JH, Johnson T, McElreath R, Smirnov O (2007) Egalitarian motives in humans. Nature 446(7137):794–796.
2 Rawls J (1971) A Theory of Justice (Harvard Univ Press, Cambridge, MA).
3 Smith A (1759) The Theory of Moral Sentiments (Millar, Kincaid, and Bell, London).
4 Darwin C (1871) The Descent of Man, and Selection in Relation to Sex (John Murray, London).
5 Rand DG, Tarnita CE, Ohtsuki H, Nowak MA (2013) Evolution of fairness in the one-shot anonymous Ultimatum Game. Proc Natl Acad Sci USA 110:2581–2586.
6 Güth W, Schmittberger R, Schwarze B (1982) An experimental analysis of ultimatum bargaining. J Econ Behav Organ 3(4):367–388.
7 Camerer CF (2003) Behavioral Game Theory: Experiments in Strategic Interaction (Princeton Univ Press, Princeton, NJ).
8 Henrich J, et al. (2006) Costly punishment across human societies. Science 312(5781):1767–1770.
9 Nowak MA, Page KM, Sigmund K (2000) Fairness versus reason in the ultimatum game. Science 289(5485):1773–1775.
10 Johnson DD, Fowler JH (2011) The evolution of overcon dence. Nature 477(7364):317–320.
11 Alexander M, Christakis NA (2008) Bias and asymmetric loss in expert forecasts: A study of physician prognostic behavior with respect to patient survival. J Health Econ 27(4):1095–1108.
12 Manapat ML, Rand DG, Pawlowitsch C, Nowak MA (2012) Stochastic evolutionary dynamics resolve the Traveler's Dilemma. J Theor Biol 303:119–127.
13 Rand DG, Nowak MA (2012) Evolutionary dynamics in nite populations can explain the full range of cooperative behaviors observed in the centipede game. J Theor Biol 300:212–221.
14 Antal T, Traulsen A, Ohtsuki H, Tarnita CE, Nowak MA (2009) Mutation-selection equilibrium in games with multiple strategies. J Theor Biol 258(4):614–622.
15 Nowak MA, Tarnita CE, Wilson EO (2010) The evolution of eusociality. Nature 466(7310):1057–1062.
16 Abbot P, et al. (2011) Inclusive tness theory and eusociality. Nature 471(7339):E1–E4, author reply E9–E10.
17 Bush RR, Mosteller F (1955) Stochastic Models for Learning, Wiley Publications in Statistics (John Wiley & Sons, New York).
18 Fowler JH (2006) Habitual voting and behavioral turnout. J Polit 68(2):335–344.
19 Wallace BR, Cesarini D, Lichtenstein P, Johannesson M (2007) Heritability of ultimatum game responder behavior. Proc Natl Acad Sci USA 104(40):15631–15634.
20 Dawes CT, et al. (2012) Neural basis of egalitarian behavior. Proc Natl Acad Sci USA 109(17):6479–6483.