SUPPORTING INFORMATION FOR J.B. ANDRÉ AND N. BAUMARD "SOCIAL OPPORTUNITIES AND THE EVOLUTION OF FAIRNESS"

JEAN-BAPTISTE ANDRÉ AND NICOLAS BAUMARD

1. THE ULTIMATUM GAME

1.1. **Nash equilibria.** As mentioned briefly in the main text, there is an infinite number of Nash equilibrium strategies (i.e. strategies that are best reply to themselves) in the ultimatum game (UG). In fact, any strategy with $p = q \in [0, 1]$ is a Nash equilibrium. If one's partner requests x (respectively if she offers x), there is nothing better to do than offering x (respectively, requesting x). Evolutionarily speaking, when every individual in the population offers and requests exactly $x \in [0, 1]$, then any mutant is either neutral or counter-selected when playing in front of the resident.

Polymorphic states of the population can also be stable. A polymorphic Nash equilibrium is a state of the population such that (i) every strategy present in the population obtains the same payoff G, and (ii) every absent strategy obtains $G' \leq G$ when confronted to a representative sample of the *resident* population. Note that, in the remaining of the Supporting Information, we sometimes use the expression "Nash equilibrium" or even sometimes only "equilibrium", for both polymorphic Nash equilibria and Nash equilibria *stricto sensu*.

In the UG, polymorphic Nash equilibria are characterized by the following properties: (i) the population is fixed for a single offer $p \in [0, 1]$ and (ii) it contains a diversity of requests $q_j \in [0, p]$ (a specific quantitative condition is also necessary on the distribution of requests, for instance, the request $q = p$ must at least be present

1

and in sufficient frequency). In any population with the above properties, infinitely rare mutants are strictly counter-selected or neutral. In particular, mutants differing with regard to their offer are all strictly counter-selected, whereas mutants differing in terms of their request are all neutral, as long as their request remains lower than the fixed offer p.

1.2. **Likely evolutionary end-points.** The Nash equilibrium condition in the UG is thus degenerate: any share of the resource is possible. But all Nash equilibria are not equally likely end-points of the evolutionary process.

First of all, one's request q is neutral in a population in which every individual offers $p \geq q$ because rejection is then never actually expressed. But rejecting any offer, even the lowest, is always maladaptive when it does occur (in game theoretic terms, strategies with $q > 0$ are Nash equilibria but they are not *subgame perfect*). Therefore, as soon as there is some variability in the population in the offers made (e.g. owing to mutation or errors), then natural selection favors the lowest possible requests. Another way to look at this issue is to remark that, in front of the resident p, all requests $q \leq p$ are perfectly neutral. Therefore, what selects among various requests is only their ability to handle non-residents (i.e. individuals with offers $p' \neq p$), and in this regard the lowest request is always the best.

However, again, this is not all that simple. Mutations (or errors) should not only affect offers, they should also affect individuals' requests. Because request is under weak selection (selection on request is exclusively dependent on the existence of some polymorphism of the offers), mutation and errors are likely to play a great role in the evolution of requests.

Therefore, the mutation rate has two counterbalancing effects. On one hand, mutations increase the strength of selection toward a reduction of requests because they generate some polymorphism on offers. But, on the other hand, mutations also introduce some noise in the requests, and thus reduce the importance of selection.

In result, what simulations and mathematical analyses show is that the issue of the evolutionary process depends precisely on the respective frequency (and effects) of mutations on individuals' offers and requests (Gale et al., 1995).

However, there is some truth in the simple intuitive reasoning. Proposers do benefit from a strategic advantage because they are the first to commit to a partition of the resource, and responders are doing the best of a bad job in accepting anything given to them. In result, when the mutation rates (on offers and requests) are equal and sufficiently large, requests and offers always evolve toward the minimum possible value (see Fig. 1 of main text).

2. Partner change and social mobility

In a polymorphic population, some strategies are more rapidly paired than others (because they offer more or request less). In result, the effective frequency of a given strategy among available partners is not equal to its actual frequency in the overall population. This should be taken into account to describe mathematically the evolutionary dynamics of a polymorphic population. In the following analyses, however, we do not aim to describe such dynamics. One just needs to keep in mind that, when we speak of the frequency or the presence of a given strategy, we mean its *effective* frequency or presence among available partners.

The simple argument developed in 4.1 of main text describes the key force at work in the model. However, like in the UG, this reasoning can be slightly misleading. In reality, individuals' request is often neutral, and selection depends on the existence of a background variability of offers, which is not necessarily present, nor extensive. There might hence exist some equilibrium states in which requests are not optimal, because the variability of offers necessary to favor an optimal request is absent. In this section of the Supporting Information, we derive the necessary properties of any

equilibrium state of the population without the *a priori* assumption that requests

2 are optimized by selection, and we show that our major result essentially holds.

2.1. **Summary.** We first summarize the major principles of the analysis.

4   First, an equilibrium state of the population must be a (polymorphic) Nash equi-

librium, i.e. (i) all the strategies present in this state (resident strategies) must

6 reach the same payoff G, and (ii) any non-resident strategy must reach a payoff

$G' \leq G$ when confronting a representative sample of the *resident* population. Here,

8 we prove that there are three families of Nash equilibria, covering all the possible so-

cial outcomes, from $p = 0$ to $p = 1$. Therefore, like in the UG, the Nash equilibrium

10 condition is degenerate, as any division of the resource can be reached.

Yet, by virtue of the second ESS condition (Maynard Smith and Price, 1973), we

12 also show that the Nash equilibria with $p < \frac{\delta}{2}$ or $p > 1 - \frac{\delta}{2}$ are not evolutionarily

stable, and not even neutrally stable (sensu Maynard Smith, 1982; see Weibull,

14 1997, p. 46) because they can be invaded by mutants that are neutral (or quasi-

neutral) when they play against the resident, but strictly favored when they play

16 "in front of themselves" (i.e. in front of other mutants like themselves). Only the

equilibria characterized by a fixed offer $p \in [\frac{\delta}{2}, 1 - \frac{\delta}{2}]$ (and a neutral diversity of

18 requests $q \leq p$, with a sufficient proportion of individuals requesting exactly $q = p$)

are neutrally stable. When the cost of postponing an interaction is low (δ close to

20 1), all the neutrally stable states of the population thus lie within a small interval

of offers around $1/2$ (precisely between $\delta/2$ and $1 - \delta/2$). In these states, all offers

22 are accepted and the average payoff is $G = 1/2$.

On top of being one of the possible neutrally stable states, the very equilibrium

24 with $p = \delta/2$ has a supplementary property related to our simple argument (section

4.1 of main text). In all neutrally stable equilibria, the average payoff being $G = 1/2$,

26 everyone expects to gain exactly $\delta/2$ in the next round, and it is thus maladaptive to

refuse offers above $\delta/2$. Yet, in neutrally stable equilibria with a fixed $p > \delta/2$, some

individuals must be ready to reject offers just below p, which *would* be maladaptive
if it were to occur. Therefore, if background polymorphism on offers is present (e.g.
owing to mutation), then the very equilibrium with $p = \delta/2$ is likely to be eventually
reached by evolution, as the best strategy in presence of polymorphism is indeed to
reject offers below $\delta/2$ but accept all others.

In what follows we explain these results in more detail.

2.2. **Nash equilibria.** Here, we consider the case in which individuals change part-
ner after each rejection and in which each partner's role is attributed at pure chance
(see sections 2 and 4 of main text for details). We first aim to characterize the
necessary properties of any (polymorphic) Nash equilibrium in this game. Recall
that a polymorphic Nash equilibrium is characterized by the following properties:
(i) all the strategies present in this state (resident strategies) reach the same payoff
G, and (ii) any non-resident strategy reaches a payoff $G' \leq G$ when confronting a
representative sample of the resident population.

In a resident population in a Nash equilibrium all the strategies present obtain
a payoff G, and therefore every individual can expect to obtain δG in his next
encounter. In what follows, we thus draw a classification of all the possible Nash
equilibria, in function of the order relationship between δG and the resident offers.

2.2.1.   The first type of potential equilibria are states of the population in which
the maximal offer present, p_{max}, is lower or equal to δG. This type of states can be
Nash equilibria only if (i) p_{max} is exactly equal to δG, and (ii) it is in fact the single
offer present. This is shown as follows.

First, consider the case in which $p_{max} < \delta G$. Because every individual expects
to gain δG in the next interaction, they are better off refusing every offer. More
formally, for every individual requesting $q_j \leq p_{max}$, one can construct a mutant with
strictly larger payoff, with the same offer but with a request $q'_j > p_{max}$. Therefore, at

a Nash equilibrium with $p_{max} < \delta G$, all requests in the population must be strictly

larger than p_{max}. In result, interactions are always rejected and the expected payoff

is $G = 0$. This leads to an inconsistency, as every offer cannot be strictly lower than

0. Therefore, a population cannot be at a Nash equilibrium with every offer strictly

lower than δG.

Consider now the case in which $p_{max} = \delta G$. Because every individual expects

to gain δG in the next interaction, they are better off refusing every offer except

p_{max}. Therefore, at an equilibrium with $p_{max} = \delta G$, all the offers below p_{max} must

be rejected. In result, if there existed some individuals offering less than p_{max}, their

payoff would be given by $G = \delta G$, which leads to an inconsistency (recall that $\delta < 1$).

Therefore, the population can be in a Nash equilibrium only if $p = p_{max} = \delta G$ is

the single offer present.

In this case, because every individual expects to gain δG in the next interaction,

acceptance or rejection of $p = \delta G$ is neutral. Assume that a fraction x of individuals

accept the offer p and a fraction $1 - x$ refuse it. The payoff of accepting individuals

writes $G_a = \frac{1}{2}p + \frac{1}{2}[x(1-p) + (1-x)\delta G_a]$, which gives $G_a = \frac{p+x(1-p)}{2-\delta(1-x)}$. Whereas the

payoff of the individuals who reject p writes $G_r = \delta \frac{G_r}{2} + \frac{1}{2}[x(1-p) + (1-x)\delta G_r]$,

which gives $G_r = \frac{x(1-p)}{2-\delta(2-x)}$. This can be in equilibrium only if $G_a = G_r$, and thus if

the effective frequency of accepting individuals is equal to $\hat{x} = \frac{2p(1-\delta)}{(1-2p)\delta}$ (this entails

that the actual frequency of accepting individuals is $\hat{h} = (\hat{x}+1)/2$). This is possible

only if $\hat{x} < 1 \Leftrightarrow p \leq \delta/2$. In this case, we necessarily have also $\hat{x} > 0$. In this

equilibrium, we verify that the average payoff of individuals, G, is indeed such that

$p = \delta G$. In the particular case in which the fixed offer is exactly $p = \delta/2$, the

equilibrium frequency of accepting individuals is $\hat{x} = 1$, i.e. every individual accepts

the offered p.

Overall, this family of equilibria corresponds in fact to two sub-cases. (i) There

is a single offer $p = \delta/2$ fixed in the population, every individual accepts it, and the

average payoff per interaction is $G = 1/2$. (ii) There is a single offer $p < \delta/2$ fixed in the population, a fraction $\hat{x} < 1$ of individuals accept it, and the average payoff per interaction is $G = p/\delta$.

2.2.2.   The second type of potential equilibria are states of the population in which there exists at least an offer strictly below δG and an offer strictly above. Call p_l the largest offer strictly below δG and p_h the smallest offer strictly above δG. This family of states can be Nash equilibria under the following *necessary* conditions.

1. *All requests fall within the interval $[p_l, p_h]$.* Because every individual expects to gain δG in the next interaction, they are better off refusing every offer below δG and accepting every offer above. Formally, for each individual requesting $q_j \notin [p_l, p_h]$, one can construct a mutant with strictly larger payoff, with the same offer but with a request $q'_j \in [p_l, p_h]$. In result, at equilibrium, all requests in the population must fall within the interval $[p_l, p_h]$.

2. *p_h is the largest offer present in the population.* Because every offer $p_i \geq p_h$ is always accepted, for each individual offering more than p_h, one can construct a strictly favored mutant with the same request but with the offer $p = p_h$. Therefore, at equilibrium, the population contains (i) "high offer" individuals all offering the same $p_h > \delta G$ which is always accepted, and (ii) a diversity of "low offer" individuals offering $p_i < \delta G$ that are always rejected.

3. *The high offer is $p_h > 1 - \delta/2$.* Call y the effective frequency of "high offer" individuals, and derive the stability condition on y. The expected payoff of "high offer" individuals can be derived as $G_h = \frac{1 - p_h(1-y)}{2 - \delta(1-y)}$, whereas the expected payoff of "low offer" individuals is $G_l = \frac{yp_h}{2 - \delta(2-y)}$. The population is at an equilibrium only if $G_h = G_l$, which yields $\hat{y} = \frac{2(1-p_h)(1-\delta)}{\delta(2p_h - 1)}$. At this equilibrium, both "high offer" and "low offer" individuals are effectively present only if $\hat{y} \in]0, 1[$, which implies $p_h > 1 - \delta/2$. Note that, in this equilibrium, every individual obtains the same payoff $G = \frac{1-p_h}{\delta}$.

4. *The population is fixed for a single request $q = p_h$.* At equilibrium, individuals with intermediate offers must be counter-selected or neutral, when they are very rare. Consider an individual offering $p' \in]p_l, p_h[$, and assume that his offer is accepted with probability x. When it is infinitely rare, this individual obtains a payoff $G' = \frac{\hat{y}p_h + x(1-p')}{2 - \delta(2 - x - \hat{y})}$, whereas residents all obtain $G = \frac{1 - p_h}{\delta}$. Simple algebra show that the rare mutant increases in frequency whenever $x > 0$, i.e. whenever it has a strictly positive probability of seeing his offer accepted. Therefore, the population is at an equilibrium only if every offer below p_h is always rejected, i.e. if it is fixed for a single request $q = p_h$. In this case, all offers strictly below p_h (including p_l) are strictly neutral.

2.2.3. The third type of potential equilibria are states of the population in which the minimal offer present, p_{min}, is strictly larger than δG. Such states can be Nash equilibria under the following *necessary* conditions:

1. *The largest request present in the population is equal to p_{min}.* This is shown as follows. First, if $p_{min} > \delta G$, every individual is better off accepting every offer. Formally, for every individual requesting $q_j > p_{min}$, one can construct a mutant with strictly larger payoff, with the same offer but with a request $q'_j \leq p_{min}$. Therefore, at an equilibrium in which $p_{min} > \delta G$, the largest request is $q_{max} \leq p_{min}$. Second, at an equilibrium in which the minimal offer is p_{min}, all mutants offering less than p_{min} must be neutral or counter-selected. Therefore, there must be at least a strictly positive frequency of individuals with $q_{max} = p_{min}$. Conclusion: the largest request present in the population is exactly equal to the smallest offer.

2. *The minimal offer p_{min} is strictly larger than $\delta/2$.* At an equilibrium in which $p_{min} > \delta G$, we have shown that we must have $q_{max} = p_{min}$.

Therefore, offers are never rejected at such an equilibrium. In consequence, the average payoff individuals gain per social interaction is $G = 1/2$ (because rejections never occur), and the smallest offer present in the population must be $p_{min} = q_{max} > \delta/2$.

3. *The population is fixed for a single offer $p > \delta/2$.* For each individual offering $p_i > q_{max}$, one can construct a mutant with strictly larger payoff, with the same request but with an offer $p_i' = q_{max}$. Therefore, in a population in equilibrium with $p_{min} > \delta G$, all offers must be exactly equal to $p = q_{max} > \delta/2$.

4. *The fixed offer is lower than $1 - \delta/2$.* This is shown as follows. At equilibrium, mutants offering less than $p = q_{max}$ must be counter-selected or neutral. Consider a mutant offering $p' < q_{max}$, such that a fraction $x < 1$ of the population accepts his offer. His payoff is given by $G' = \frac{1}{2}p + \frac{1}{2}x(1 - p') + \frac{1}{2}(1 - x)\delta G'$, which yields $G' = \frac{p + x(1 - p')}{2 - \delta(1 - x)}$. This payoff depends on x, i.e. on the actual composition of the resident population in terms of request. But G' can potentially be lower than $1/2$ at least for certain values of x only if it is lower than $1/2$ when $x = 0$. This yields the simple necessary condition $p \leq 1 - \delta/2$. Said differently, if $p > 1 - \delta/2$, then mutants offering less than p are favored even if they are certain to see their offer rejected. Therefore, a population at an equilibrium with $p_{min} > \delta G$ is necessarily fixed for a single offer $p \leq 1 - \delta/2$.

5. We also check that mutants requesting more than $p = q_{max}$ are counter-selected or neutral. The payoffs of these mutants is given by $G' = \frac{1}{2}(1 - p) + \frac{1}{2}\delta G'$, which yields $G' = (1 - p)/(2 - \delta)$. Simple algebra shows that this payoff is lower than the residents' payoff ($1/2$) iff $p \geq \delta/2$. By definition, this is always true as we have already shown that a population with $p_{min} > \delta G$ can be an equilibrium only if $p_{min} > \delta/2$.

2.2.4.   There is a fourth type of potential equilibria in which there exists at least an offer $p_h > \delta G$ that every responder accepts, and an offer $p_l = \delta G$ that an intermediate fraction x of responders accept. However, here there is no mechanism that leads the average payoff G to be equal to p_l/δ. Therefore, this family of equilibrium is extremely unlikely to be reached exactly.

Let us now recapitulate. There are three types of (polymorphic) Nash equilibria in this game, covering all the possible social outcomes.

1. A single offer $p < \delta/2$ is fixed that a fraction $\hat{x}$ of responders accept. The average payoff per individual is then $G = p/\delta$ (see section 2.2.1).

   The intuitive logic behind these equilibria is the following. When offers are low in the population, the role of proposers is more desirable than the role of responders. Therefore, when one happens to play the role of a responder in a given interaction, it may pay to reject the offer, and postpone the interaction until the next encounter, in which one may have a chance to play the more favorable role of proposer. But the more responders reject offers, the lower is the payoff one can expect to gain as a proposer. At one point it does not even pay anymore to reject offers. Therefore, the proportion of individuals who reject p reaches an intermediate equilibrium. This equilibrium is precisely characterized by the fact that the expected payoff of each responder in the next interaction (would he reject the current offer), δG, is exactly equal to the offer p.

2. A single request $q > 1 - \delta/2$ is fixed, and a fraction $\hat{y}$ of individuals offer exactly $p = q$, whereas a fraction $1 - \hat{y}$ offer less. The average payoff per individual is then $G = (1 - q)/\delta$ (see section 2.2.2).

   The intuitive logic behind these equilibria is the symmetric of the first. When requests are high in the population, the role of responders is more desirable, and it may pay, as a proposer, to offer less than requested, in order to be rejected and have a chance to play the role of responder in the

next encounter. The proportion of individuals who fulfill the requested q is subject to negative frequency dependence and stabilizes when the expected payoff of each proposer in the next interaction (would he not fulfill the fixed request), δG, is exactly equal to the effective "offer" made $1 - q$.

3. A single offer $p \in [\frac{\delta}{2}, 1 - \frac{\delta}{2}]$ is fixed, and it is always accepted (individuals request a variety of $q_i \leq p$). The average payoff per individual is exactly $G = 1/2$ (see sections 2.2.1 and 2.2.3).

The intuitive logic behind this third category of equilibria is in line with the simple argument developed in main text (section 4.1). Because every individual has a fair chance of playing the role of a responder or the role of a proposer, and because rejections never occur, every individual can expect to gain exactly $\delta/2$ in the *next* interaction. In result, as long as individuals gain more than $\delta/2$ in the *current* interaction, they should be happy with it (accept the offer, or comply to the request). This entails that the resource division is not too strongly biased in favor of one role or the other (each role must obtain at least $\delta/2$).

In the following, it will be useful to separate these states in two sub-types: (i) states in which the fixed offer is in the interval $]\frac{\delta}{2}, 1 - \frac{\delta}{2}]$ (section 2.2.3), and (ii) states in which the fixed offer is exactly $p = \delta/2$ (section 2.2.1).

2.3. **Neutral stability.** Like in the UG, the Nash equilibrium condition is degenerate in this model, as any division of the resource can be reached. However, we must now check that these equilibria remain stable when polymorphism is present. Recall that, when a resident population is at a Nash equilibrium, many mutants are perfectly neutral (in particular mutants differing with regard to their request), and the fate of these mutants depends entirely on their ability to handle "non-resident" strategies present as background polymorphism. Here, we show that this phenomenon is particularly prone to destabilize the two first types of Nash equilibria, and

can destabilize all but one equilibrium state of the third type, but to a lesser extent.

2 Consider the three cases in turn.

2.3.1.   Let us consider the first type of equilibria (section 2.2.1). Consider a pop-

4 ulation containing a single offer $p < \delta/2$ and two types of requests, a fraction $\hat{x}$ of

accepting individuals with $q \leq p$, and a fraction $1 - \hat{x}$ of rejecting individuals with

6 $q > p$. Recall that, by definition, at equilibrium, rejecting and accepting individuals

are in an equilibrium frequency in which they both receive the same expected payoff.

8 These states are potentially Nash equilibria in the sense that there exists some pre-

cise combinations of requests among the rejecting individuals such that all mutants

10 are neutral or counter-selected. But the necessary conditions on this composition is

very stringent and the equilibrium is in fact very easily destabilized.

12 On one hand, if there exists a sufficient density of (neutral) rejecting individuals

whose request q' is larger but very close to p, then mutants offering $p' \gtrsim q'$ can

14 be favored because the small cost of their generosity is compensated by the fact

that their offer is accepted more often than the resident's. On the other hand, and

16 conversely, the fact that the population contains some generous mutants offering $p' >$

p, favors rejecting individuals who request $q' \in]p, p']$, because (i) rejection is neutral

18 in a purely resident population offering p (by definition of the equilibrium), and (ii)

strictly advantageous in presence of some generous mutants because it increases the

20 likelihood of ending up paired with such a mutant.

These equilibria are thus destabilized by the co-evolution of large offers and large

22 requests. In equilibrium, both large offers and large requests are either neutral or

counter-selected when taken individually, but any initial introduction of one of the

24 two mutants favors the other and vice versa, which destabilizes the equilibrium.

More formally, in what follows we show that this process can occur in particular

26 when the very same mutants both offer and demand more than p. Consider a

resident population with two strategies: (i) accepting individuals, in frequency $\hat{x}$,

2 offering $p < \delta/2$ and requesting $q \le p$, and (ii) rejecting individuals, in frequency

$1 - \hat{x} - \epsilon$, offering the very same $p < \delta/2$ but requesting $q > p$. The remaining of

4 the population consists of mutants, in effective frequency ϵ, offering and demanding

$p' = p + \delta$. For simplicity, we assume that all rejecting residents have $q > p'$

6 (this is the worst case scenario for the mutants). The payoff of accepting residents

writes $G_a = \frac{\hat{x}(1-p)+p+\epsilon\delta}{2-\delta(1-\hat{x})}$; the payoff of rejecting residents writes $G_r = \frac{\hat{x}(1-p)}{2-\delta(2-\hat{x})}$,

8 and the payoff of mutants $G' = \frac{\hat{x}(1-p-\delta)+\epsilon}{2-\delta(2-x-2\epsilon)}$ (with $\hat{x} = \frac{2p(1-\delta)}{\delta(1-2p)}$). To first order in

δ, the payoff difference between accepting and rejecting residents is $\frac{(1-2p)\epsilon\delta}{2(1-p)-\delta} \ge 0$

10 and therefore, the accepting individuals have the largest payoff of all residents. The

payoff difference between mutants and accepting residents is then given by $\Delta_a =$

12 $A\epsilon - B\delta + o(\delta)$, with $A = \frac{(1-2p)^2\delta}{C} > 0$, and $B = \left[\frac{2p(1-\delta)}{C} + \frac{(1-2p)}{2(1-p)-\delta}\epsilon\right] > 0$ (where

$C = 2\delta[(1-\delta)(1-p)+\epsilon\delta(1-2p)] > 0$). Therefore, for any initial mutant's frequency,

14 ϵ, one can always find mutants with a sufficiently weak effect δ, that are able to rise

from this initial frequency (i.e. $\Delta_a > 0$). This occurs because the mutants are

16 perfectly neutral in a population of pure residents, but strictly favored in presence

of some mutants like themselves (see also Fudenberg and Maskin, 1990; Binmore

18 and Samuelson, 1992; André and Day, 2007).

2.3.2. Let us now consider the second type of equilibria (section 2.2.2). Consider a

20 population containing a single request $q > 1 - \delta/2$ and two types of offers, a fraction

$\hat{y}$ of "high offers" ($p = q$), and a fraction $1 - \hat{y}$ of "low offers" ($p_i < q$). Recall that,

22 by definition, at equilibrium, "high offer" and "low offer" individuals both receive the

same expected payoff.

24 These states are also easily destabilized by the co-evolution of offers and requests.

Here, low requests ($q' < q$) and low offers ($p' < q$) are both strictly neutral in equi-

26 librium, but they favor each other. The presence of few neutral mutants requesting

less than q makes it advantageous to offer less than q, and vice versa.

More formally, we can show that this process occurs in particular when the very same mutants both offer and demand less than q. Consider a resident population with two strategies: (i) "high offer" individuals, in frequency $\hat{y}$, requesting $q > 1 - \delta/2$ and offering $p = q$, and (ii) "low offer" individuals, in frequency $1 - \hat{y} - \epsilon$, requesting the very same $q < 1 - \delta/2$ but offering $p' < q$. The remaining of the population consists of mutants, in effective frequency ϵ, requesting and offering $q' = q - \delta$. For simplicity, we assume that all "low offer" residents have $p < q'$ (this is the worst case scenario for the mutants). In equilibrium, when $y = \hat{y}$ mutants are rare, all residents obtain the same payoff; and the difference between mutants are residents payoff is strictly equal to $\frac{(1-2q)^2\epsilon}{2[q(1-\delta)+(2q-1)\epsilon\delta]} > 0$, which is independent on δ. Therefore, mutants rise in frequency.

Terminologically speaking, the two first types of equilibria are polymorphic Nash equilibria, but they are not evolutionarily stable, and not even neutrally stable (sensu Maynard Smith, 1982's and Weibull, 1997, p. 46), because "neutral" (or "quasi-neutral") mutants can increase in frequency by creating for themselves a favorable environment.

2.3.3. Let us now consider the third type of equilibria (section 2.2.3). Consider a population containing a single offer $p \in]\frac{\delta}{2}, 1 - \frac{\delta}{2}]$ that every individual accepts (i.e. all requests are lower than p). Here, we show that they are also prone to destabilization by mutant-mutant interaction, but to a lesser extent.

These states are equilibria because any deviation is neutral or counter-selected. In particular, all responder's strategy are neutral because they never actually reject anything. But things are different in the presence of variability on offers. If responders are regularly confronted to deviants who offer less than p, then rejection is not neutral anymore. In fact, a responder can expect to gain exactly $\delta/2$ in the next social interaction (if she rejects her current offer). Therefore, responders are always better off if they accept any offer larger than $\delta/2$ in the first place.

More precisely, consider a responder requesting exactly $q = p$, and compare it with another responder requesting exactly $q = \delta/2$. These two strategies end up with different payoffs, and the strategy with $q = \delta/2$ fares strictly better, when they are confronted to an individual offering $p \in [\delta/2, p[$. Therefore, in the presence of some background polymorphism on offers falling precisely within the interval $[\delta/2, p[$, the Nash equilibrium with $p > \delta/2$ becomes unstable because lower requests are favored.

However, the mutants whose offer fall within the interval $[\delta/2, p[$ are not directly favored by selection, as the mere presence of some mutants requesting $\delta/2$ is not sufficient to favor them. They become favored by selection only once the average requests has significantly decreased below p. And the average request can significantly decrease only if mutants offering less than p are present. Therefore, here, the destabilization of the equilibrium relies on the recurrent introduced of counter-selected variability owing to mutation or other sources of variation.

Terminologically speaking, this third type of equilibria are (polymorphic) Nash equilibria, and they are also neutrally stable in the sense that no rare mutant can increase in frequency (Maynard Smith and Price, 1973). But they are not strict ESS and, more significantly, they are not perfect equilibria in the sense that they do not stipulate a best reply in front of all deviants (Selten, 1975, 1983, 1988; even though the evolutionary translation of the perfection concept is not straightforward, as it depends on the nature of variability; see Gale et al., 1995; Binmore and Samuelson, 1999).

However, here as well, terminological issues are not central. More significantly, whereas the two first types of states are destabilized by combinations of *advantaged* mutants creating a favorable environment for each other, this third type of states are destabilized by the ongoing presence of *counter-selected* mutants (with lower p) who create a favorable environment for otherwise neutral mutants (with lower q).

This second process thus relies on some variability introduced recurrently in the population.

2.3.4.   In contrast with all other equilibria, the states in which the population is fixed precisely for the offer $p = \delta/2$ and contains a polymorphism of requests $q_i \leq p$ is immune to destabilization by all mutant-mutant interactions. On one hand, mutants that differ with regard to their offer are strictly counter-selected when they are confronted to the resident population. On the other hand, mutants that differ with regard to their request are at best neutral both in front of the resident and in front of any possible deviant because, whatever is the offer one is confronted to, playing the request $q = \delta/2$ is the best possible strategy when one expects to gain exactly $\delta/2$ in the next interaction.

Terminologically speaking, the single strategy $(\frac{\delta}{2}, \frac{\delta}{2})$ is a Nash equilibrium, it is neutrally stable (Maynard Smith and Price, 1973), and it is also a perfect equilibrium, in the sense that it stipulates a best reply even in front of deviants who do not follow the same strategy (Selten, 1975, 1983, 1988). But it not a strict ESS because many mutants (with lower requests) are perfectly neutral when this strategy is entirely fixed.

3. PARTNER SWITCHING WITHOUT SOCIAL MOBILITY

Assume that individuals are assigned a given role at birth (either proposer or responder) that they will always play in every social interaction. Assume that the frequency of each role is controlled by extrinsic mechanisms (e.g. the occurrence of deleterious mutations) and cannot evolve in response to the payoff obtained in each role.

To understand the outcome of the model under this assumption, we follow the same simple argument as in 4.1 of main text) under the assumption that responders'

request is optimized by selection. We show that proposers obtain the entire resource
at evolutionary equilibrium. The argument develops as follows:

1. In the population at equilibrium, there must exist a payoff G_r that all re-
    sponders gain in average in each social interaction.

2. Natural selection favors responders who request exactly $q = \delta G_r$, i.e. re-
    sponders always request a little bit less than their average payoff.

3. Natural selection favors proposers who offer exactly $p = q = \delta G_r$ (because
    proposers who offer less obtain a nil payoff, and proposers who offer more
    are unnecessarily giving away resources).

4. Therefore, responders obtain $p = \delta G_r$ in each interaction, and their expected
    payoff G_r must satisfy the condition $G_r = \delta G_r$.

5. As long as $\delta \neq 1$, $G_r = 0$ is the only possible payoff for responders at an
    evolutionary equilibrium: the fixed offer is $p = 0$, and the request $q = 0$.

4. Social mobility without partner switching

Assume that, when an offer is rejected by a responder, the two partners remain
together and their respective role in the interaction is simply re-attributed at ran-
dom. Partners are separated only once an actual interaction has taken place (i.e.
once an offer has been accepted), except if the interaction turns out to be impossible
(because the partners' offers and requests are incompatible), in which case the two
individuals are separated (with no payoff).

Here we develop a very simple argument to show that the only *monomorphic*
equilibrium in this case occurs when the population is fixed with an offer p and a
request $q = p$, with $p \in [\delta/2, 1 - \delta/2]$. The argument develops in two steps.

1. *The population is fixed with responders requesting q and proposers offering
    exactly $p = q$.* Assume that the population is fixed with a strategy (p, q)
    with $p < q$. The average payoff per social interaction is $G = 0$ because

offers are always rejected. A mutant proposer offering $p' = q$ obtains a payoff $G' = (1 - q)/2 > G$ and thus increases in frequency, thereby showing that (p, q) with $p < q$ cannot be stable. On the contrary, assume that the population is fixed with a strategy (p, q) with $p > q$. The average payoff per social interaction is $G = 1/2$ because offers are always accepted. A mutant proposer offering $p' = q$ obtains a payoff $G' = (1 + p - q)/2 > 1/2$ (because $p > q$), and thus increases in frequency, thereby showing that (p, q) with $p > q$ cannot be stable. If a monomorphic population is to be stable, it must be fixed with a strategy (p, q) with $p = q$.

2. *All individuals offer and request $p \in [\delta/2, 1 - \delta/2]$ at equilibrium.* When the population is fixed with $p = q$, the average payoff per social interaction is $G = 1/2$. Assume that the population is fixed with $p < \delta/2$. In this case, a mutant requesting $q' > p$ obtains a payoff $G' = \frac{1-p}{2-\delta}$, which is strictly larger than $1/2$ (because $p < \delta/2$). Therefore the mutant is strictly favored and the population is not in equilibrium. Assume that the population is fixed with $p > 1 - \delta/2$. In this case, a mutant offering $p' < p$ obtains a payoff $G' = \frac{p}{2-\delta}$, which is also strictly larger than $1/2$ (because $p > 1 - \delta/2$). Therefore the mutant is strictly favored and the population is not in equilibrium either. When the population is fixed with $p \in [\delta/2, 1 - \delta/2]$, in contrast, all mutants are either neutral or counter-selected.

5. Simulation procedure

We performed stochastic individual-based simulations. The simulation program, coded in C, is available upon request. Simulations work as follows.

The population size is fixed. Generations are non-overlapping. Individuals are haploids. They are genetically characterized by (i) their offer when they play the role of a proposer and (iii) their minimal request when they play the role of a responder.

Mutations occur at the same rate at each locus. There is no recombination between
loci.

The life cycle is as follows. In each generation, (1) social interactions take place
and individuals accrue a social payoff, (2) a non social payoff is added, (3) mutation
takes place, (4) the next generation is sampled as a multinomial distribution in
which the expected number of adult offspring of each individual is equal to its payoff
divided by the total payoff in the population.

The social procedure (step 1) has four versions depending on the model.

1. Ultimatum game. All individuals are randomly paired; a proposer is ran-
domly chosen in each pair; the interactions take place (or not) and the pairs
are separated. Individuals hence have a single *opportunity* of social interac-
tion per generation.

2. Partner switching and social mobility. In this second model, we aim to con-
sider the fact that individuals can reject a current opportunity of social in-
teraction in order to reach a later (potentially better) opportunity (probably
with a different partner). In our analytical model, we assume that indi-
viduals can possibly undertake a fixed number of actual interactions across
their life. Therefore, each accepted opportunity entails the implicit renunci-
ation to another, and this generates a selective pressure to accept only good
enough opportunities. In our simulations we consider the following, more
general, situation. The social life of individuals lasts for an explicit number
L of time steps. Across these L time steps, individuals can be in two states:
a state S in which they can interact socially (they are "susceptible" to social
interactions) and a state R in which they can't (they are "resistant" to social
interactions). They move from S to R each time they do entertain a social
interaction, and they move from R to S with a constant probability ρ in

each time step. This means that, after each accepted social interaction, individuals undergo a "refractory" period of expected length $1/\rho$, during which they are non-receptive to further interaction. $1/\rho$ could be interpreted, for instance, as the actual duration of each social interaction. We assume that at the beginning of their life individuals are in the state R (but similar results are obtained in the opposite case, not shown). In each time step, the individuals in state S are randomly paired (one of them remains single if the number of S individuals is uneven), a proposer is chosen in each pair, and interactions take place, or not, depending on offers and requests in each pair. Partners are then separated, moved to state R if necessary, and so on. When an offer is accepted, payoffs are accrued to each individual of the pair. These payoffs are multiplied by a factor δ^t for each individual, in which $\delta < 1$ is the discount factor, and t is the number of time steps that have past since the individual last moved from state R to state S.

3. Partner switching without social mobility. The third model is exactly identical to the second, except that each individual's role in a pair of partner is not assigned randomly but in function of intrinsic individuals' properties. At birth, every individual is randomly assigned a role (proposer or responder, with equal probabilities) that she keeps for all her life, and will pay in every social interaction she may undertake.

4. Social mobility without partner switching. The fourth model is similar to the first one. Individuals undertake only one social interaction in their life, with a single partner. However, the social interaction is more complex than in the ultimatum game. It explicitly consists of a number L of time steps. At each time step, one of the partners in each pair is chosen as the proposer, makes an offer that the other accepts or rejects. When an offer is rejected, the interaction moves on to the next time step. When an offer is accepted, payoffs are given to each partner and the interaction is over. The payoff of each individual is multiplied by δ^t, in which δ is the discount factor, and

t is the number of time steps that have past since the beginning of the interaction. After L time steps, all interactions are terminated no matter what, and the individuals reproduce.

6. Simulation results

Here we present some complementary simulation results in the version 2 of the model, i.e when partner switching and social mobility co-occur.

First, in the main text, we presented only the average offer in this model (Fig. 2 of main text). He, we show both the average offer and the average request (Fig. 1) with the same parameter set (and with initial conditions $p = q = 0.1$). The average request follows the same pattern as the average offer with faster stochastic variations due to the fact that requests are often neutral.

Second, in the main text we presented simulation results with generation length $L = 10^3$ and $\rho = 10^{-2}$. Under this assumption, after each interaction, individuals undergo in expectation a period of 100 time steps during which they cannot interact further. In such circumstance, the cost of being picky is moderate because the total number of actual interactions an individual can undergo is strongly limited anyway. This leads to outcome that are very close to our analytical predictions (Fig. 2 of main text). In the sake of comparison, here we present results in different conditions. We keep the generation length $L = 10^3$ constant, but we consider different values of ρ. These results are shown in Figure 2. When ρ is larger, the cost of being picky is stronger because the total number of actual interactions an individual can undergo is less limited. Individuals are thus better off accepting even relatively low offers. This leads to a division of the resource that is more favorable to proposers. In the extreme, when $\rho = 1$, individuals can immediately interact again after each social interaction, therefore they should never reject any offer, and the model thus

corresponds to the ultimatum game (proposers obtain the entire resource at equilibrium; dotted line in Fig. 2). This illustrates the fact that our results do depend strongly on the assumption that there exists a *competition* between various social opportunities, in such a way that rejecting an opportunity opens another one later.

REFERENCES

André, J. and T. Day, 2007. Perfect reciprocity is the only evolutionarily stable strategy in the continuous iterated prisoner's dilemma. J Theor Biol 247:11–22.

Binmore, K. and L. Samuelson, 1992. Evolutionary stability in repeated games played by finite automata. Journal of Economic Theory 57:278–305.

———, 1999. Evolutionnary drift and equilibrium selection. Review of Economic Studies 66:363–393.

Fudenberg, D. and E. Maskin, 1990. Evolution and cooperation in noisy repeated games. The American Economic Review 80:274–279.

Gale, J., K. Binmore, and L. Samuelson, 1995. Learning to be imperfect: The ultimatum game*. Games and Economic Behavior 8:56–90.

Maynard Smith, J., 1982. Evolution and the Theory of Games. Cambridge University Press, Cambridge.

Maynard Smith, J. and G. R. Price, 1973. The logic of animal conflict. Nature 246:15–18.

Selten, R., 1975. Reexamination of the perfectness concept for equilibrium points in extensive games. International Journal of Game Theory 4:25–55.

———, 1983. Evolutionary stability in extensive two-person games. Mathematical Social Sciences 5:269–363.

———, 1988. Evolutionary stability in extensive two-person games - correction and further development. Mathematical Social Sciences 16:223–266.

Weibull, J., 1997. Evolutionary Game Theory. MIT Press, Cambridge MA.
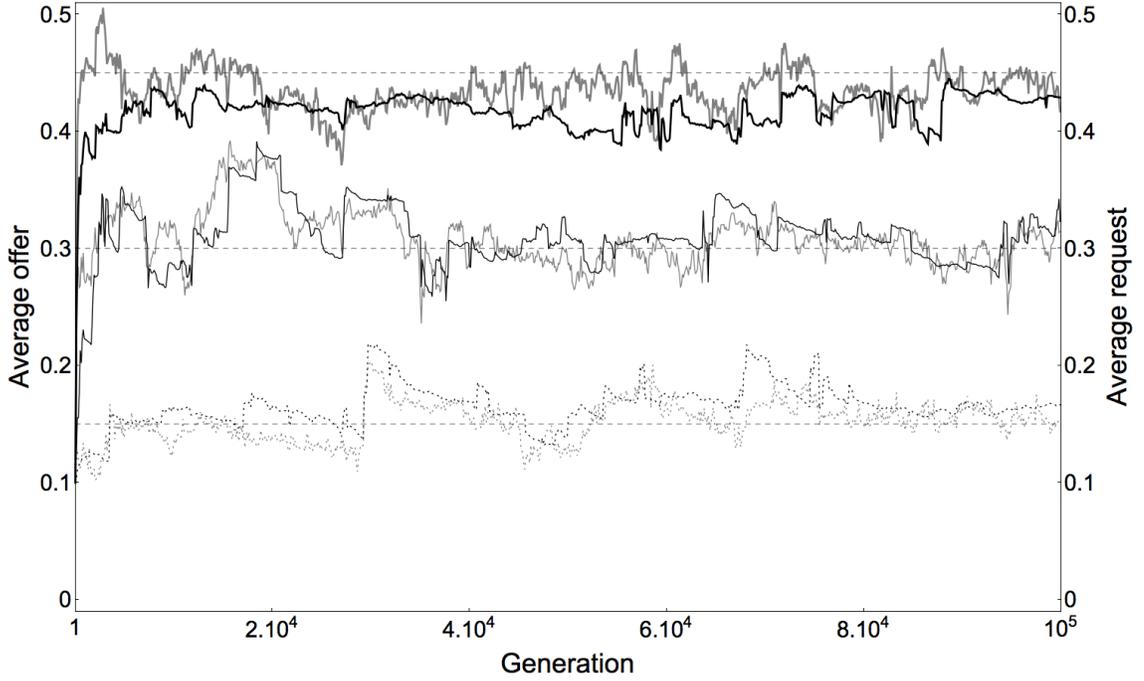
FIGURE 1. *Offers and requests.* We simulate numerically the evo-
lution of a population of individuals playing the UG when they can
change partner and change role (see section 5). Each curve is an
average over 10 simulation runs. The simulations are initiated with
fixed offer and request $p = q = 0.1$. We consider three discount
factors, $\delta = 0.9$ (thick plain lines), $\delta = 0.6$ (thin plain lines), and
$\delta = 0.3$ (dashed lines). Average offers are shown in black and aver-
age requests in grey. The straight dashed lines show the respective
analytical predictions with each discount factor. All parameters are
like in Fig. 2 of main text: $L = 10^3$, $\rho = 0.01$, non-social payoff is
10^{-5}, population size $N = 10^3$, mutation probability $\mu = 10^{-3}$, a
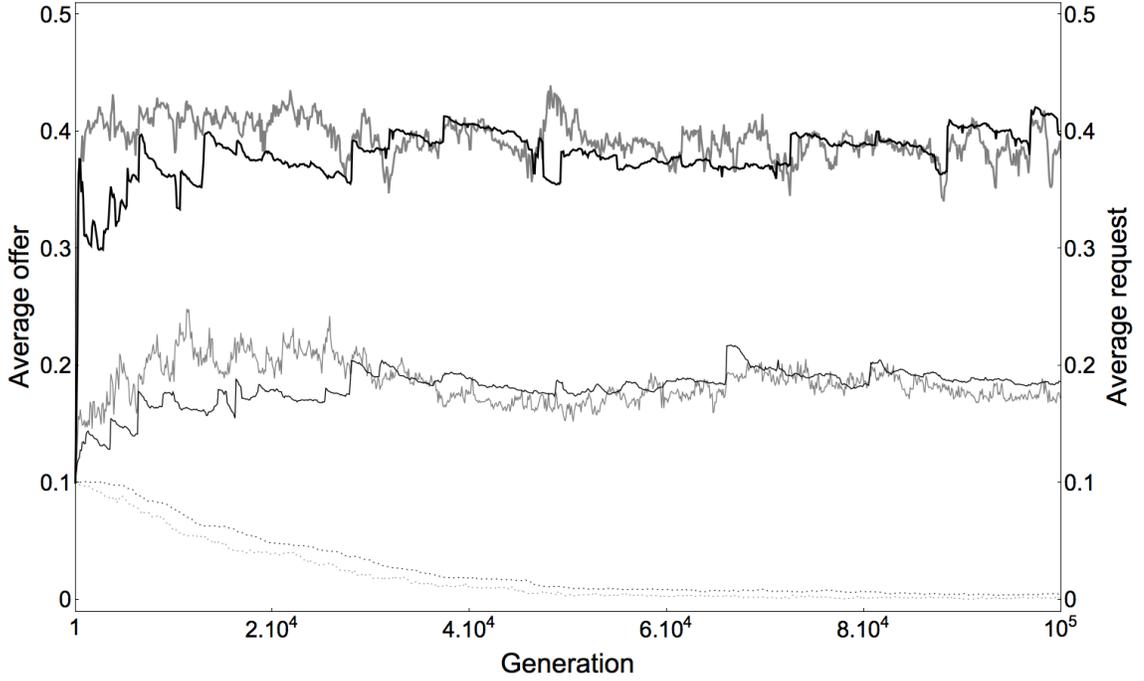fraction 0.1 of which have a strong effect.

FIGURE 2. *Faster rate of social interaction.* We simulate numerically the evolution of a population of individuals playing the UG when they can change partner and change role (see section 5). Each curve is an average over 10 simulation runs. The simulations are initiated with fixed offer and request $p = q = 0.1$, the discount factor is $\delta = 0.9$. Average offers are shown in black and average requests in grey. All parameters are like in Fig. 1, except $\rho = 0.1$ (thick lines), $\rho = 0.5$ (thin lines), and $\rho = 1$ (dotted lines).