

Changing behaviour under unfairness: An evolutionary model of the Ultimatum Game

Arioli Gianni, Lucchetti Roberto, Valente Giovanni *

Politecnico di Milano, Department of Mathematics, Milan, Italy

ARTICLE INFO

Keywords:

Fairness
Rational behaviour
Evolutionary game theory
Dynamical evolution

ABSTRACT

Experimental results on the Ultimatum Game indicate that receivers may reject non-zero offers, even though that seems irrational. The explanation is that, when players are treated unfairly, they can act against strict rationality. This paper discusses an evolutionary model of the Ultimatum Game describing how populations of players change their behaviour in time. We prove an analytical result that establishes under what conditions receivers tend to reject unfair offers. The response to unfair offers is also shown to be sensitive to different degrees of unfairness. We then introduce a Bayesian game to translate our result from populations to individual players.

1. Rationality and (un-)fairness

In the Ultimatum Game two players are supposed to share a certain amount of money, say 100 dollars. One player, i.e. the proposer, is entitled to make an offer and the other player, i.e. the receiver, can either accept or reject it: if she accepts the offer, then both players keep the share they agree upon; else, if she rejects the offer, then neither player gains anything. From a strictly rational point of view, in order to maximize her utilities, the receiver should accept any non-zero offer, even when it corresponds to a very low share, for in this way she would at least gain something rather than nothing. Nevertheless, as Güth et al. [5] and Güth [9] demonstrated,¹ experiments conducted in the laboratory show that receivers tend to reject offers that they deem too low. In fact, data indicate that in most cases an agreement between the proposer and the receiver is reached at a 65 : 35 ratio, while less fair ratios such as 80 : 20 would be typically rejected. This poses two related philosophical issues. To begin with, it appears that standard game theory fails against empirical tests. That raises the question whether there is any alternative way to give a game-theoretical description of the observed behaviour. Moreover, one is left with the problem of explaining why real-life bargaining deviates from full rationality. Arguably, the purported explanation is that people's actions are subject to social and cultural norms, which determine certain behavioural patterns. One should thus identify the normative conditions that are factored into the relevant decision processes and build a model that combines them with rational constraints, in such a way to replicate the observed behaviour. The present paper takes up these outstanding issues, by developing a dynamical model of the Ultimatum Game within the framework of evolutionary game theory.

The idea underlying our proposal is that behavioural norms have an evolutionary underpinning. In fact, the mutual interaction between agents evolves according to a kind of stimulus-response mechanism, whereby certain patterns of behaviour survive while

* Corresponding author.

E-mail address: giovanni.valente@polimi.it (G. Valente).

¹ See also Güth and Yaari [7], Güth and Schmittberger [8], Güth and Kocher [6].

<https://doi.org/10.1016/j.ijar.2023.109078>

Received 22 August 2023; Received in revised form 25 October 2023; Accepted 28 October 2023

Available online 8 November 2023

0888-613X/© 2023 The Author(s). Published by Elsevier Inc. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).




Fig. 1. The game tree for Ultimatum Minigame.

others are abandoned. Evolutionary game theory provides a framework to identify long-run patterns, as they arise from dynamical equations describing the behaviour of different populations of players interacting in specific bargaining scenarios. In the case of the Ultimatum Game, the strategies adopted by proposers and receivers are sensitive to the norm of *fairness*. For, even when the game is played anonymously, the receiver would tend to refuse low offers made by the proposer as long as she perceives them to be too unfair. The desire to be treated in a fair manner thus explains why the pattern of behaviour observed in the experiments survives notwithstanding the fact that refusing a non-zero offer does not seem fully rational (cfr. Aumann [1,2]; Debove et al. [3]; Huck and Oechssler [11]; Marchetti et al. [12]; Nowak et al. [14]). In other words, fairness places a normative constraint into the decision process, over and above the standard maximization of utilities. An evolutionary model of the Ultimatum Game can show how the interaction between the players varies in the course of time until a long-run equilibrium is reached, thereby selecting determinate patterns of behaviour. The model we elaborate here draws from a numerical example presented by Gale, Binmore and Samuelson [4, sec.5] as a variant of Selten's [16] Chain Store Paradox. While the original aim of these authors was to study interactive learning processes under perturbations, in our formulation of the replicator equations we keep the values of the offers being made variable. This puts us in a position to evaluate how the response changes with respect to different "degrees of unfairness". In addition, by applying the population model to the realistic scenario in which the game is played by individual agents, it enables us to confront its results with the experimental behaviour observed for the Ultimatum Game. For this purpose, it is insightful to translate our evolutionary model into a Bayesian game featuring two rational players, who interact with each other according to the rules of the Ultimatum Game, but are supposed to take decisions in face of uncertainty. The deterministic dynamics dictated by the replicator equations is thus replaced by individual probabilistic reasoning. Specifically, the Bayesian game we construct aims to establish what the rational response of the receiver should be when the proponent makes unfair offers. As we shall see, though, the choice of the appropriate probability function encoding the information that results from the numerical simulations of our evolutionary model is far from trivial.

The paper is structured as follows. In section 2, we construct our evolutionary model wherein the mutual interaction between two distinct but equinumerous populations of proposers and receivers is dictated by the respective replicator equations, by taking into account how the relative utilities of the players evolve through time. It is worth stressing that it yields a genuinely asymmetric version of the Ultimatum Game, which depends on both the values of the fair and the unfair offers the proposers would make to the receivers. The analysis we develop is grounded on a proposition that we prove, which establishes the populations' behaviour in the infinite-time limit for any possible initial condition. Based on our formal result, we then survey the outcomes of actual simulations of the model for different values of the unfair offer. In the following section 3, we show how to translate the result obtained for populations into a probabilistic framework for individual decisions. We can therefore determine under what conditions it is more convenient for a receiver playing the Ultimatum Game to accept or reject an unfair offer. In particular, it turns out that the observed behaviour is a function of the "degrees of unfairness" of the offer the proposer makes to receiver. Finally, in section 4, we add some conclusive remarks on the significance of the dynamical model discussed here as an alternative game-theoretical account for the behaviour observed in experiments on the Ultimatum Game.

2. An evolutionary model

We consider and further elaborate a version of the Ultimatum Minigame that has been introduced in Section 5 of Gale et al. [4], represented by the tree in Fig. 1.

Accordingly, both players have two strategies: Player P can choose between a fair offer a or an unfair offer b . If the offer is fair, the game ends; otherwise while Player R can choose whether to accept or reject the unfair offer. The choice of the parameters a and b is arbitrary, at least so long as one takes $a > b \geq 0$, in the sense that a fair offer is by definition higher than an unfair offer, which in turn cannot be negative. To fix the idea, as a standard value for a fair offer one may take $a = 1/2$, namely the perfectly equal share among the players,² whereas the offer b can vary below such a value at Player P's will, thereby defining a continuous range of degrees of unfairness. One can in fact expect a different response from Player R depending on how much unfair the offer made by Player P is perceived: for instance, the low offer $b = 1/10$ may appear so unfair that Player R would tend to reject it, as the experiments on the Ultimatum Game indicate. It is exactly for the purpose of evaluating the behaviour of the players on the basis of distinct degrees of unfairness that we introduce variables for the fair offer a and the unfair offer b . Let us emphasize that the

² Although this seems the most natural choice for the fair offer, it should be noted that there is a long-standing discussion in the philosophical literature on game theory regarding whether fairness should correspond to equal share among all players. See Skyrms [17] for a justification of this fact, especially in the context of the game known as "divide the cake".

general setting we propose here differs from Gale et al.'s [4] original proposal exactly on this point, in that they roughly categorise the offers into low and high ones, instead of sharply distinguishing between fair and unfair offers as we do. By letting the unfair offer vary below the fair offer we thus wish to offer a quantitative analysis of the extent to which low offers become too unfair to be acceptable. As we will demonstrate in the course of the paper, the exact relationship between the parameters a and b proves crucial for the behaviour of the model. For any allowed choice of a and b , we can represent the game in strategic form by the bimatrix:

$$A = \begin{pmatrix} (1-a, a) & (1-a, a) \\ (1-b, b) & (0, 0) \end{pmatrix} \quad (1)$$

This is the standard formulation of the Ultimatum Minigame. Specifically, the first column corresponds to the case in which Player R accepts both the fair offer (first line) and the unfair offer (second line). Instead, the second column corresponds to the case in which Player R refuses the unfair offer: hence, if Player P makes the fair offer (first line), both players gain their respective share, but if Player P makes the unfair offer (second line), both players end up with nothing.

The model then considers two distinct equinumerous populations P and R , one comprising Proposers and the other comprising Receivers, respectively. It is posited that every member of a population meets regularly a random member of the other population to play the Ultimatum Game in the above formulation. The P population is thus split into a fraction x of all members who play fair and a fraction $(1-x)$ of members who play unfair; instead, the R population is split into a fraction y of members who play accept everything and a fraction $(1-y)$ of members who play reject the unfair offer.

Next, we proceed to determine the utilities for the different strategies the players can adopt, both on the individual level for the members of each population and on average with respect to the whole population. The utility is computed by the amount of money a random player in one population will gain for a given offer, weighted by a factor corresponding to the fraction of members of the other population that goes along with that offer. Accordingly, the utility for a member of the P population is $(1-a)$ if she plays fair, since all receivers will accept the offer a and hence the weight is equal to 1; whereas the utility is $y(1-b)$ if she plays unfair, since only a fraction y of receivers will accept the offer b . Likewise, the utility for a member of the R population is $xa + (1-x)b$ if he plays accept, since a fraction x of the proposers will make a fair offer a and the remaining fraction $1-x$ will make an unfair offer b ; whereas, the utility is xa if he plays reject the unfair offer, since he will only gain the fair offer a weighted by the fraction x of proposers who make it. In order to compute the average utility, for each population we sum over the individual utilities multiplied by the corresponding fraction of members. It follows that, for the P population, the average utility is $x(1-a) + (1-x)y(1-b)$. Instead, for the R population, it is $yxa + y(1-x)b + (1-y)xa$, which after a straightforward calculation becomes $xa + y(1-x)b$. According to evolutionary game theory, the size of the sub-populations x and y will change through time depending on the relative values of the thus-defined utilities.

In the dynamical approach, the evolution of each sub-population is dictated by replicator equations. Such equations define the relative rate of change of a given sub-population as the difference between the individual utility of its members and its average utility.³ So, based on the utility values we have computed above, we can form the following set of replicator equations for the unknown variables x and y :

$$\begin{cases} \dot{x} = (1-a) - [x(1-a) + (1-x)y(1-b)], \\ \dot{y} = xa + (1-x)b - [xa + y(1-x)b]. \end{cases} \quad (2)$$

The right-hand sides of these equations can be simplified to:

$$\begin{cases} \dot{x} = [(1-a) - y(1-b)] \cdot x(1-x), \\ \dot{y} = b(1-x) \cdot y(1-y). \end{cases} \quad (3)$$

Notice that the lines $x=0$, $x=1$, $y=0$, $y=1$ are invariant, and therefore the square $[0, 1] \times [0, 1]$ is also invariant. Moreover, since one has $\dot{x}=0$ if $x=0$ or $x=1$ and one has $\dot{y}=0$ if $y=0$ or $y=1$, within the boundaries of such a unit square for any possible initial conditions $(x_0, y_0) \in (0, 1)^2$, there exists a solution $(x(t), y(t)) \in (0, 1)^2$ of (3) for all time t . The time-evolution of the model depends on the initial conditions, that is the fraction $x_0 = x(0)$ of proposers who make a fair offer and the fraction $y_0 = y(0)$ of receivers who accept any offer at the initial time $t=0$. The relevant quantities of interest are then given by the values the variables x and y take on when time grows to infinity, namely $x_\infty := \lim_{t \rightarrow +\infty} x(t)$ and $y_\infty := \lim_{t \rightarrow +\infty} y(t)$. In order to see how the two sub-populations evolve towards these limit values, we need to study each equation in the system (3) in greater detail.

According to the second replicator equation, the fraction y of the R population of receivers that accept all offers, including an unfair offer, varies with respect to three non-negative terms: a non-negative factor that depends on its own size, i.e. $y(1-y)$, the value

³ Let us stress that our approach is purely deterministic in that it is based on replicator equations, and as such it differs from some recent attempts to cast the Ultimatum Game in evolutionary terms, such as that by Nowak et al. [15]. In fact, while the latter also consider two distinct finite populations of proponents and receivers, they construct a stochastic dynamics for fairness whereby random fluctuations can occur with a certain probability. In particular, they argue that such randomness may facilitate increase of fairness as a form of self-interested natural selection. Instead, our approach does not introduce any random component and simply aims at establishing how mutual interactions between the players vary in the course of time under deterministic evolution equations as a result of different degrees of unfairness. Yet another interesting approach is found in the work by Zisis et al. [18], which accounts for the evolution of fairness (i.e. the level of generosity) of the proponent. However, this approach does not quite refer to the Ultimatum Game, but rather it discusses the Dictator Game as a primary example of anticipation games: as such, despite containing valuable insights, it is less relevant for the matter addressed in the present paper.

of the unfair offer b and the fraction $1 - x$ of the P population that make the unfair offer. As such, the rate of change \dot{y} is greater than zero, which means that $y(t)$ increases in the course of time. In fact, one always has $\dot{y} > 0$, except in the case when $1 - x = 0$: so, unless there is no member of the sub-population x that make an unfair offer, the number of receivers y that accept all offers keeps growing, and its rate of growth increases as the number of proposers x that make the fair offer gets lower. What is more, for $(x, y) \in (0, 1)^2$ one has $\dot{y} \rightarrow 0$ if $x \rightarrow 1$, which is consistent with the fact that it is always more convenient for a receiver to accept all offers, but the advantage decreases when a larger part of the proposers plays fair. Note also that, for each fixed pair of x and y at a given time, the rate of growth of y becomes higher as the value of the unfair offer b raises: in other words, a larger fraction of the R population would tend to accept even an unfair offer if the latter is, so to speak, less unfair. We will return to this point in the next section when discussing the notion of probability of acceptance for different degrees of unfairness. For now, let us focus on the limit value of the variable y . Provably, since for all initial conditions $(x_0, y_0) \in (0, 1)^2$ one has $\dot{y} = by(1 - x)(1 - y) > 0$ and hence the solution $y(t) \in (0, 1)$ is always strictly increasing, but since it is bounded it admits a finite limit y_∞ as $t \rightarrow +\infty$. In particular, it follows that $y_\infty := \lim_{t \rightarrow +\infty} y(t) \leq 1$. Accordingly, it could as well happen that not all receivers will eventually accept any offer, that is there could be a non-zero fraction $1 - y_\infty$ of the R population that will still reject the unfair offer. Here below we prove an analytical result establishing the circumstances under which the limit value y_∞ becomes 1, and thus all the members of the R population accept any offer. Yet, before presenting our full result, we ought to discuss how the behaviour of the P population is supposed to change in the course of time.

According to the first replicator equation, the fraction x of the P population of proposers that make the fair offer varies with respect to a non-negative factor that depends on its own size, i.e. $x(1 - x)$, as well as with respect to another term $[(1 - a) - y(1 - b)]$ that depends on the value of the fair offer a , the value of the unfair offer b and the fraction y of the R population that accept all offers. This latter term vanishes when y is equal to $\frac{1-a}{1-b}$, which in turn implies that the number of proposers that play fair remains stable, i.e. $\dot{x} = 0$. The threshold $\hat{y} := \frac{1-a}{1-b}$ thus determines the sign of \dot{x} : that is, if $y < \hat{y}$ the fraction of proposers x that play fair tends to increase, i.e. $\dot{x} > 0$, whereas if $y > \hat{y}$ the fraction x tends to decrease, i.e. $\dot{x} < 0$, even though the fraction y of the R population that accept all offers still grows. The existence of such a threshold is consistent with the fact that it is more convenient for a proposer to play unfair just in case a sufficiently large portion of receivers will accept it. As one can readily see, by settling the fair offer $a = 1/2$, the value of \hat{y} increases for higher values of the unfair offer b , meaning that the threshold grows based on the extent to which the offer made by the proposers is perceived by the receivers as being less unfair (we will compare the outcomes of numerical simulations of the model for different choices of b at the end of the section). What is perhaps more surprising, though, is that, once the fair offer a and the unfair offer b are fixed, the behaviour of the P population in the limit for time going to infinity crucially depends on the bound \hat{y} , as the following Proposition shows.

Proposition 1. *Consider the system*

$$\begin{cases} \dot{x} = [(1 - a) - y(1 - b)] \cdot x(1 - x), \\ \dot{y} = b(1 - x) \cdot y(1 - y). \end{cases} \quad (4)$$

Let $a > b > 0$, and let $\hat{y} := \frac{1-a}{1-b}$. Then, for all initial conditions $(x_0, y_0) \in (0, 1)^2$ we have $\lim_{t \rightarrow +\infty} (x(t), y(t)) = (x_\infty, y_\infty) \in \{0, 1\} \times [0, 1]$. Furthermore, there exists an increasing function $f : [0, 1] \rightarrow [0, \hat{y}]$ such that $f(0) = 0$, $f(1) = \hat{y}$ and

- For all initial conditions $(x_0, y_0) \in (0, 1)^2$ such that $y_0 > f(x_0)$ we have $x_\infty = 0$ and $y_\infty = 1$.
- For all initial conditions $(x_0, y_0) \in (0, 1)^2$ such that $y_0 \leq f(x_0)$ we have $x_\infty = 1$ and $y_\infty \in (0, \hat{y}]$.

Proof. First note that for all initial conditions $(x_0, y_0) \in (0, 1)^2$ the replicator equation admits a solution $(x, y) : \mathbb{R} \rightarrow (0, 1)^2$, and $y(t)$ is strictly increasing. Let $y_{-\infty} = \lim_{t \rightarrow -\infty} y(t)$ and $y_\infty = \lim_{t \rightarrow +\infty} y(t)$. If $y_\infty > \hat{y}$, then $y(t) > \hat{y}$ for all t larger than some \bar{t} (possibly $\bar{t} = 0$), and then x is strictly decreasing for $t > \bar{t}$, and therefore $x(t) \rightarrow x_\infty$ as $t \rightarrow \infty$. Then

$$\lim_{t \rightarrow +\infty} \dot{x} = \lim_{t \rightarrow +\infty} x(1 - x)(1 - a - y(1 - b)) = x_\infty(1 - x_\infty)(1 - a - y_\infty(1 - b)),$$

therefore $x_\infty(1 - x_\infty)(1 - a - y_\infty(1 - b)) = 0$, but since $y_\infty > \hat{y}$ and $x(t)$ is decreasing, then $x_\infty = 0$. If instead $y_\infty \leq \hat{y}$, then $y(t) < \hat{y}$ for all t , therefore x is always increasing, and because of a same argument as before $x(t) \rightarrow x_\infty = 1$. Similarly, one can prove that $x_{-\infty} = \lim_{t \rightarrow -\infty} x(t) = 0$ and $y_{-\infty} = 0$ for all initial conditions $(x_0, y_0) \in (0, 1)^2$.

Now let $(x_c(t), y_c(t))$ be the solution of the initial value problem $(x_0, y_0) = (1/2, c)$, with $c \in (0, 1)$. Clearly, if $c > \hat{y}$, then $(x_c(t), y_c(t)) \rightarrow (0, 1)$. On the other hand, if c is sufficiently close to 0, then $(x_c(t), y_c(t)) \rightarrow (1, y_\infty)$, with $y_\infty \leq \hat{y}$. To prove this last statement, we need to show that, for small c , it is $y_\infty < \hat{y}$. Consider the trajectories in the rectangle $R_k = [1/2, 1] \times [0, \hat{y}/2]$. In this region $\dot{x} > 0$, therefore $x(t)$ is invertible, and if $Y(x) = y(t(x))$ we have

$$Y'(x) = \frac{\dot{y}}{\dot{x}} = \frac{bY(1 - Y)}{x(1 - b)(\hat{y} - Y)} \leq \frac{4Y}{\hat{y}},$$

so that the solution of

$$Y'(x) = \frac{bY(1 - Y)}{x(1 - b)(\hat{y} - Y)}, \quad Y(1/2) = c$$

satisfies

$$Y(x) \leq ce^{\frac{2}{y}}.$$

It follows that, if $c > 0$ is sufficiently small, the trajectory through $(\frac{1}{2}, c)$ is such that $y(t) \leq \hat{y}/2$ for all $t \geq 0$, and therefore $(x_c(t), y_c(t)) \rightarrow (1, y_\infty)$.

Let $\bar{c} = \sup c$, where the sup is taken among all the values of c such that $x_c(t) \rightarrow 1$. Then $\lim_{t \rightarrow \infty} (x_{\bar{c}}(t), y_{\bar{c}}(t)) = (1, \hat{y})$. Furthermore, both $x_{\bar{c}}(t)$ and $y_{\bar{c}}(t)$ are strictly increasing, so $x_{\bar{c}}(t)$ is invertible. Let $t_{\bar{c}} : (0, 1) \rightarrow \mathbb{R}$ be the inverse, and define $f : [0, 1] \rightarrow [0, \hat{y}]$ by $f(x) := y_{\bar{c}}(t_{\bar{c}}(x))$. Then f has the required properties, and since trajectories cannot intersect because the system is autonomous, all trajectories starting on or below the graph of f converge to $(1, y_\infty)$. On the other hand, all trajectories starting strictly above the graph of f intersect the line $x = 1/2$ at some point $(1/2, y)$ with $y > f(1/2)$, therefore they converge to $(0, 1)$ because of the definition of \bar{c} . \square

Note that Proposition 1 assures that the limit value $x_\infty := \lim_{t \rightarrow +\infty} x(t)$ is either 0 or 1, that is all the proposers will eventually play unfair or play fair, respectively, depending on whether or not the fraction of receivers that accept all offers eventually exceeds the threshold \hat{y} . Furthermore, Proposition 1 establishes the circumstances under which $y_\infty = 1$, which can only happen when $x_\infty = 0$.

This analytic result puts us in a position to evaluate the limit behaviour of the sub-population x of proposers who play fair given any initial value (x_0, y_0) . If $y_0 > f(x_0)$, then it is always the case that $y_\infty = 1$, and therefore $x_\infty = 0$. It means that, if there is a sufficiently large fraction of the R population that accepts any offer at the initial time, the members of the P population that make an unfair offer tend to increase to a point that eventually no one will play fair anymore, while all members of the R population will end up accepting the unfair offer. If $y_0 \leq f(x_0)$, instead, it may happen that at the beginning $x(t)$ increases, but eventually $y(t)$ becomes larger than \hat{y} , and at that point $x(t)$ begins decreasing, and continues to decrease until it vanishes. Then $y(t)$ is forced to increase until it reaches 1 (in an infinite time). It means that, if there is only a small fraction of the R population that accepts the unfair offer at the initial time, the members of the P population will tend to play fair. Note that in this case, differently from the previous case, one has $y_\infty < 1$ so that not all receivers would accept the unfair offer, but that does not make any actual difference since all the offers made by the proposers are fair.

It is interesting to evaluate the behaviour of each sub-population on the basis of the different degrees of unfairness. In our model, the size of the unfair offer b can vary continuously in $(0, a)$, that is, with our choice of fair offer, in $(0, 1/2)$. Of course, the higher the value of b , the lower the degree of unfairness of the unfair offer, in the sense that the latter would get closer to the fair offer a . Let us stress that the evolution of x and the evolution of y depend on b in rather different ways. On the one hand, the number of receivers y in the R sub-population who accept any offer, including an unfair one, increases linearly with respect to b . Hence, \hat{y} is very small for very low offers, while it tends to grow rapidly as the unfair offer becomes less unfair. On the other hand, the evolution of the number of proponents x in the P sub-population who make a fair offer crucially depends on the function $f(x)$, which depends on the unfair offer b . Specifically, Proposition 1 shows that, if $(x(t), y(t))$ is the solution of (4) with initial conditions (x_0, y_0) , then either $\lim_{t \rightarrow +\infty} (x(t), y(t)) = (0, 1)$, that is the whole population converges to the standard solution of the Ultimatum Minigame, or $\lim_{t \rightarrow +\infty} (x(t), y(t)) = (1, y_\infty)$, that is the whole P population makes a fair offer, and then it does not matter what R does. In Fig. 2 we display (in blue) the graph of the function $f(x)$ introduced in the Proposition, together with the trajectories for some choice of initial conditions, for different values of the unfair offer b (while the fair offer remains $a = 1/2$) ordered in terms of decreasing degrees of unfairness, namely $b = 0.01, b = 0.1, b = 0.25, b = 0.35, b = 0.4, b = 0.49$. In Fig. 3 we plot the area below the graph of f , corresponding to the cases in which all x eventually play fair, as a function of b . We remark that such area approaches $1/2$ both when $b \rightarrow 0$ and when $b \rightarrow 1/2$. Indeed, in the first case the graph of $f(x)$ tends to the horizontal line $y = \hat{y} = 1/2$, while in the second case the graph of $f(x)$ tends to the diagonal of the square. This sequence reveals that the percentage of proponents who end up making a fair offer follows a kind of parabola behaviour. Indeed, it is relatively high for very unfair offers, due to the fact that for small b the growth of y is not very fast and so it is difficult to overcome the threshold \hat{y} , even though the latter is not too high; then, it progressively decreases until it reaches its lower value around $b = 0.25$ (notice, though, that the percentage does not change much within the central interval $b \in [0.1, 0.4]$); after that, it begins to rise up again as the growth of y makes it easier to overcome the threshold \hat{y} , even though the latter increases together with b too. As the last picture indicates, when the unfair offer b gets very close to the fair offer a , so that $\hat{y} = 1$ in the limit case in which $b \rightarrow a$, the P sub-population would tend to split equally between proponents who play fair and proponents who play Unfair.

3. From populations to individuals

The evolutionary version of the Ultimatum Minigame we constructed in the previous section is framed in terms of populations. More to the point, it features two equinumerous populations of proponents and receivers interacting in such a way that at each time bargaining takes place between random pairs of members of the P population and the R population, respectively. As such, the outcomes one obtains in the course of time, both in terms of the analytic result we derived and the numerical simulations we discussed, pertain to the collective behaviour of populations, given certain initial conditions about their distribution at time $t = 0$. However, the standard Ultimatum Game is a non-dynamical scenario involving just two individual players. In the present section, we wish to translate our result from populations to individuals. Arguably, it requires one to cast the collective outcomes of the evolutionary model with dynamical equations (4) into a probabilistic framework that applies to the individual behaviour of the players, without any dependence on the initial conditions of the respective sub-populations. Let us show how this can be done in a mathematically rigorous manner.

Our proposal is to use the results provided by the evolutionary version of the Ultimatum Minigame to introduce a simple Bayesian game. Bayesian games were introduced in Harsanyi [10], see also the alternative definition by Aumann in Maschler, Solan, Zamir

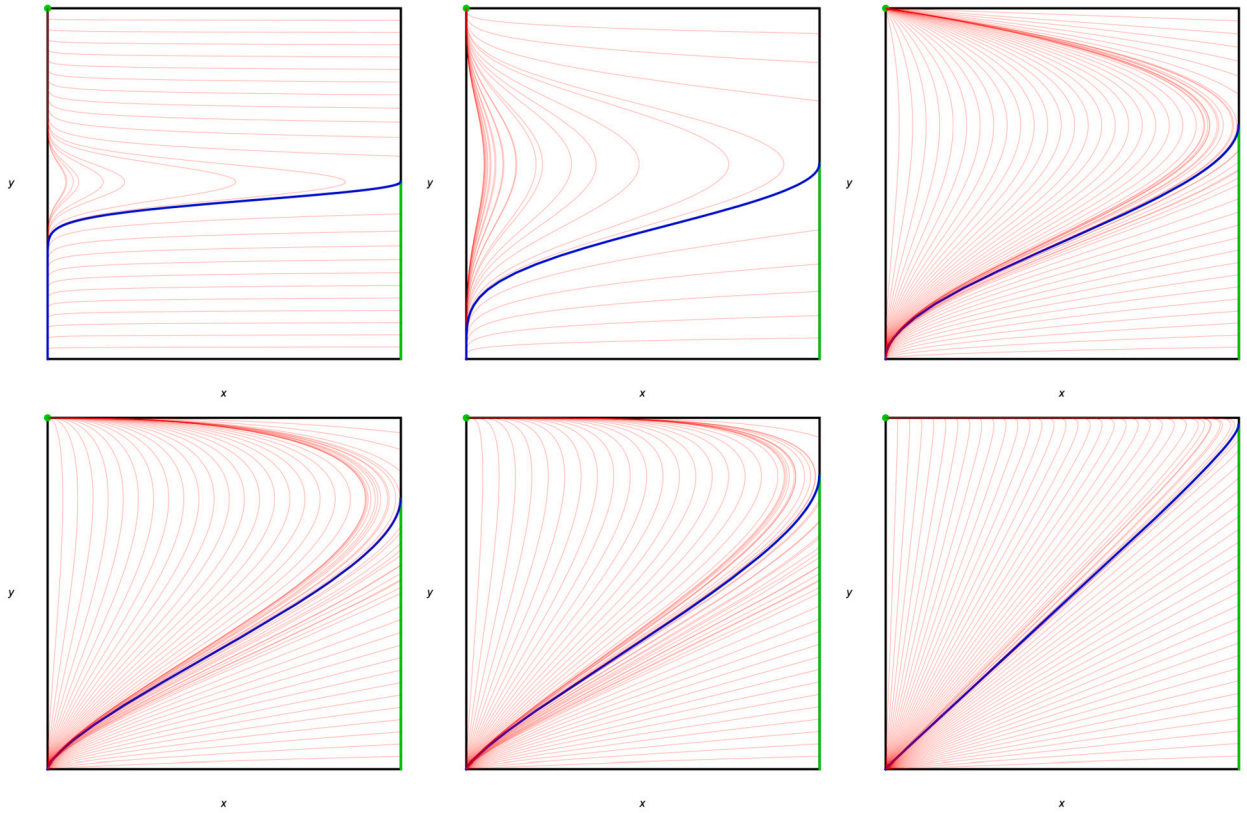


Fig. 2. Initial conditions strictly above the blue line lead to $x(t) \rightarrow 0$ (all P s eventually play unfair), while initial conditions below or on the blue line lead to $x(t) \rightarrow 1$ (all P s eventually play fair). The red curves represent some trajectories of the system. (For interpretation of the colours in the figure(s), the reader is referred to the web version of this article.)




Fig. 3. Area of the region below the blue line as a function of b .

[13], to study strategic decision-making situations where players have incomplete information. Arguably, that puts us in a position to show how the populations' behaviour predicted by the determinist replicator equations of our evolutionary model of the Ultimatum Game can be cast into a probabilistic scheme for an individual reasoning in the face of uncertainty. In fact, in our example we assume that the Proponent has only probabilistic information on the possible reaction of the Receiver to an unfair offer. Accordingly, we design a game in which there is probability p_A that player R, namely the receiver, accepts any offer from the proponent P, and probability $p_B = 1 - p_A$ that the receiver R rejects the unfair offer. This scenario can be described in extensive form as the game tree represented in Fig. 4.

Within the Bayesian framework, the proponent has one type only, with two strategies: make a fair offer or make an unfair offer. The receiver R is of two possible types with two strategies: namely, *type 1* being the player that is accepting both offers, and *type 2* being the player accepting the fair offer but rejecting the unfair offer. Here, an unfair offer b results into a negative utility for type 2, which we set to be $-\epsilon$. As it is described in the game tree of Fig. 4, Nature chooses the type of the receiver, and the proponent is unaware of the result of the choice. The proponent chooses the fair or the unfair offer. If the offer is fair, the game is finished and the outcome is $(1 - a, a)$, independently of the strategy adopted by the receiver. If the offer is unfair, the receiver of type 1 will accept anyway, since the payoff is $b > 0$, while the receiver of type 2 will reject, since the payoff is $-\epsilon < 0$.





Fig. 4. The game tree for the case of two players.

Fig. 5. p_A^1 (red, as in equation (5)) and $(1-a)/(1-b)$ (black) computed as a function of b , with $a = 1/2$.

The utilities of the proponent P in this Bayesian game are listed in the following table (expressing the strategies of the receiver R as *accept/accept*, *accept/reject*, *reject/accept*, *reject/reject*):

$$\begin{pmatrix} 1-a & 1-a & 1-a & 1-a \\ 1-b & p_A(1-b) & (1-p_A)(1-b) & 0 \end{pmatrix}$$

Provably, there are multiple Bayes-Nash equilibria, in the sense that one can opt for various possible strategies in order to maximize the utilities of each player. However, as a matter of fact, only the strategy *accept for type 1 of R and reject for type 2 of R*, which corresponds to the second column in the above table, provides a subgame perfect equilibrium. This means that the rational play for the proponent is to choose the best option between $1-a$ and $p_A(1-b)$. It follows that she is better off making a fair offer if and only if $1-a > p_A(1-b)$. In other words, player P will play fair just in case the probability p_A that player R accepts any offer is strictly smaller than the threshold $\hat{y} = \frac{1-a}{1-b}$ for the population model. By keeping $a = 1/2$ fixed, the threshold for such probabilities is therefore a function of the degrees of unfairness of the unfair offer, which depends on the variable b .

The formal connection of this fact with the analytic result we proved in the previous section is enforced by the structural definition of the probability function p_A . Indeed, since each player is supposed to be representative of the corresponding population, it is natural to think of its properties as some kind of average over the properties of the members of the population. There are two key elements in the construction of the individual probability that player R accepts whatever offer the other player P would make. For one, in the evolutionary model with dynamical equations (4), the relative portion of receivers in the population R who will eventually accept any offer is given by the asymptotic value y_∞ to which the system evolves from each possible initial condition (x_0, y_0) : as Proposition 1 shows, in some cases these values is equal to 1, but in other cases it remains below the threshold \hat{y} . The probability function p_A should thus encode the collective behaviour of the population R with respect to the limit quantity y_∞ . Furthermore, while the evolution of the populations is sensitive to their initial conditions, the behaviour of a generic player is largely independent from the distributions (x_0, y_0) at time $t = 0$: indeed, the latter can be randomly selected within all possible values in the square $(0, 1)^2$. So, when computing individual quantities, one ought to average over all possible distributions, and since there is no reason to prefer a particular one over the other, they all be assigned the same weight.

These considerations lead to a first proposal for p_A , that is

$$p_A^1(b) = \int y_\infty(x_0(b), y_0(b)) dx_0 dy_0, \quad (5)$$

where the integral is extended to the square $[0, 1]^2$. Fig. 5 shows the graph of $p_A^1(b)$, together with the graph of the function $b \mapsto \frac{1-a}{1-b}$. This choice does not match the results of the experiments: indeed, if it were the actual probability, the proponent would be better off by offering 0, since in this case her expected utility would be $p_A^1(0) \approx 0.7 > 1-a = 0.5$. Therefore, formula (5) needs to be modified.

For the sake of constructing a more suitable function expressing the probability p_A that player R accepts any offer, a few additional remarks are in order. For one, it seems reasonable to assume that p_A should be close to 0 when b approaches 0, and should be close to 1 when b approaches a , which is clearly not the case for p_A^1 . Moreover, another requirement for p_A is that for the graph of the




Fig. 6. p_A^2 (red) and $(1-a)/(1-b)$ (black) computed as a function of b , with $a = 1/2$.

function to be concave, since the probability of the receiver to accept the unfair offer should be proportional to the utility for the receiver, and the utility versus the actual gain is usually a concave function.

So, in order to formulate a better choice of the sought-after probability function p_A , we conjecture that only the initial conditions that matter are those leading to an asymptotic state where the receiver does not accept any offer. Accordingly, we could use equation (5), but extending the integral only to the region below the blue line. For this to result into a concave function, we further conjecture that such a function is raised to a suitable power $\gamma \in (0, 1)$, where we estimate γ in order to obtain the best fit with the experimental data. Thus, our proposal for p_A is the following formula:

$$p_A^2(b) = \left(\alpha \int y_\infty(x_0(b), y_0(b)) dx_0 dy_0 - \beta \right)^\gamma, \quad (6)$$

where the integral is taken over all initial conditions leading to $x_\infty = 1$ and α, β are chosen in such a way that $p_A^2(0) = 0$ and $p_A^2(a) = 1$, while the choice of γ which provides the best fit of the data is $\gamma = 0.265$. Then we obtain the result shown in Fig. 6.

As it turns out, the condition $p_A(1-b) < 1-a$ under which the proponent P is better off making the fair offer is fulfilled when $b \leq 0.34$, thereby explaining why the proponent chooses to offer at least 0.34. In this regard, it is worthwhile emphasizing that the thus-obtained result fact is fully consistent with the conclusions of Güth et al. [5] experiment on the Ultimatum Game we mentioned in the opening section, whereby an agreement between the proposer and the receiver on how to share the overall sum at their disposal is typically reached starting from a 65 : 35 ratio.

4. Conclusion

In this paper, we discussed an evolutionary model of the Ultimatum Game, wherein the replicator equations (4) for the behaviour of two (equinumerous) populations of proposers and receivers describe how the relative utilities of the players evolve in the course of time. Compared to Gale et al.'s [4] original example, the replicator equations we formulated incorporate the unfair offer as a non-zero variable b that can range continuously below the value of the fair offer $a = 1/2$, thereby putting one in a position to evaluate how the behaviour of each population changes under different degrees of unfairness. According to our model, the sub-population y of receivers who accept any offer, even an unfair one, tends to increase in the course of time, but its rate of change is slower for lower values of b , i.e. for higher degrees of unfairness. Meanwhile, so long as there are members of the R population who reject the unfair offer, there will be proposers in the P population who are better off playing fair. An analytical result we proved, namely Proposition 1 in section 2, determines the long-run behaviour of the two populations. In particular, it shows that, if the initial conditions (x_0, y_0) are such that the fraction of receivers y_∞ who eventually accept any offer remains below a given threshold \hat{y} , which grows as the offers become less unfair, then all proposers x_∞ tend to make a fair offer, since otherwise their offer can be rejected. This fact about populations was then translated into the context of a Bayesian game for individual players we constructed in section 3. Accordingly, the behaviour of an individual receiver is encoded in a probability function p_A defined by averaging the long-run distribution of population R over all possible initial conditions. We could thus show that, if p_A remains below the value of the threshold \hat{y} , it turns out to be more convenient, and therefore rational, for the proposer to make a fair offer, else the receiver would reject the offer. Interestingly, this happens for values of the unfair offer below $b = 0.34$, which is consistent with the experimental results observed in the laboratory for the Ultimatum Game.

CRedit authorship contribution statement

All authors have equally contributed to the research.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

No data was used for the research described in the article.

References

- [1] R. Aumann, Rule-rationality versus act-rationality, Discussion paper 497, 2008, pp. 1–20.
- [2] R. Aumann, A synthesis of behavioural and mainstream economics, *Nat. Hum. Behav.* 3 (2019) 666–670.
- [3] S. Debove, N. Baumard, J-B. André, Models of the evolution of fairness in the ultimatum game: a review and classification, *Evol. Hum. Behav.* 37 (2016) 245–254.
- [4] K. Gale, J. Binmore, L. Samuelson, Learning to be imperfect: the ultimatum game, *Games Econ. Behav.* 8 (1995) 56–90.
- [5] W. Güth, R. Schmittberger, B. Schwarze, An experimental analysis of ultimatum bargaining, *J. Econ. Behav. Organ.* 3 (1982) 367–388.
- [6] W. Güth, M.G. Kocher, More than thirty years of ultimatum bargaining experiments: motives, variations, and a survey of the recent literature, *J. Econ. Behav. Organ.* 108 (2014) 396–409.
- [7] W. Güth, W. Yaari, Explaining reciprocal behavior in simple strategic games: an evolutionary approach, in: U. Witt (Ed.), *Explaining Process and Change: Approaches to Evolutionary Economics*, Michigan University Press, 1992, pp. 23–34.
- [8] W. Güth, R. Schmittberger, An experimental analysis of ultimatum bargaining, *Nat. Hum. Behav.* 3 (2019) 666–670.
- [9] W. Güth, An evolutionary approach to explaining cooperative behavior by reciprocal incentives, *Int. J. Game Theory* 24 (1995) 323–344.
- [10] J.C. Harsanyi, Games with incomplete information played by “Bayesian” players, I-III part I. The basic model, *Manag. Sci.* 14 (3) (1967) 159–182.
- [11] J. Huck, S. Oechssler, The indirect evolutionary approach to explaining fair allocations, *Games Econ. Behav.* 28 (1999) 13–24.
- [12] A. Marchetti, I. Castelli, K.M. Harlé, A.G. Sanfey, Expectations and outcome: the role of proposer features in the ultimatum game, *J. Econ. Psychol.* 32 (3) (2011) 446–449.
- [13] M. Maschler, E. Solan, S. Zamir, *Game Theory*, 2nd ed., Cambridge University Press, Cambridge, 2020.
- [14] K.M. Nowak, M.A. Page, K. Sigmund, Fairness versus reason in the ultimatum game, *Science* 289 (2000) 1773–1775.
- [15] K.M. Nowak, et al., Evolution of fairness in the one-shot anonymous ultimatum game, *Proc. Natl. Acad. Sci.* 110 (7) (2013) 2581–2586.
- [16] R. Selten, The chain-store paradox, *Theory Decis.* 9 (1978) 127–159.
- [17] B. Skyrms, *Evolution of the Social Contract*, Cambridge University Press, 1996.
- [18] I. Zisis, et al., Generosity motivated by acceptance-evolutionary analysis of an anticipation game, *Sci. Rep.* 5 (1) (2015) 18076.