

Fairness in Multiplayer Ultimatum Games Through Moderate Responder Selection

Fernando P. Santos¹ and Daan Bloembergen²

¹Department of Ecology and Evolutionary Biology, Princeton University, New Jersey, USA

fpsantos@princeton.edu

²Centrum Wiskunde & Informatica, Amsterdam, The Netherlands

d.bloembergen@cwi.nl

Abstract

We study the evolution of fairness in a multiplayer version of the classical Ultimatum Game in which a group of N Proposers offers a division of resources to M Responders. In general, the proposal is rejected if the (average) proposed offer is lower than the (average) response threshold in the Responders group. A motivation for our work is the exchange of flexibilities between smart energy communities, where the surplus of one community can be offered to meet the demand of a second community. In the absence of any Responder selection criteria, the co-evolving populations of Proposers and Responders converge to a state in which proposals and acceptance thresholds are low, implying an unfair exchange that favors Proposers. To circumvent this, we test different rules which determine how Responders should be selected, contingent on their declared acceptance thresholds. We find that selecting moderate Responders optimizes overall fairness. Selecting the lowest-demanding Responders maintains unfairness, while selecting the highest-demanding individuals yields a worse outcome for all due to frequent rejected proposals. These results provide a practical message for institutional design and the proposed model allows testing policies and emergent behaviors on the intersection between social choice theory, group bargaining, competition, and fairness elicitation.

Introduction

Many social dilemmas in society can be formulated and studied using game theoretic methods (Gintis, 2000). In particular, the question how cooperation can come about in a society of self-interested individuals has attracted considerable interest in the research community (Axelrod and Hamilton, 1981; Hofmann et al., 2011; Rand and Nowak, 2013; Ranjbar-Sahraei et al., 2014; Santos et al., 2018). Typically such social dilemmas are cast as a normal form game, in which a set of players simultaneously and without prior communication choose an action to play, and the resulting joint action determines the payoff to each player. Despite the simplicity of these one-shot interactions, normal form games can still capture many of the intricate dynamics of complex strategic interactions (Axelrod and Hamilton, 1981; Skyrms, 2004).

One example of such a game is the Ultimatum Game (UG) (Güth et al., 1982), in which one player, the Proposer, offers a certain split of a resource to a Responder, who decides to accept or reject the offer. If accepted, the players receive their share per the offer; if rejected both players receive nothing. We propose and study a Multiplayer version of the classical Ultimatum Game, in which a group of N Proposers offers a division of resources to a group of M Responders. Henceforth we refer to this interaction as **NvM-person Ultimatum Game (NM-MUG)**. While a multiplayer version of UG was previously discussed in the context of one Proposer and N Responders (Santos et al., 2015, 2016, 2019), considering proposals by groups of Proposers is relevant in the context of rival public goods division, where 1) Proposers may be tempted to free-ride and lower their proposals expecting other Proposers to compensate and 2) the group sizes of Proposers and Responders may not match, reducing the per-capita share in one of the groups. In general, we assume that a proposal is rejected if the (average) proposed offer is lower than the (average) response threshold in the Responders group. We study under which conditions a fair outcome is achieved, in which Proposers offer a substantial split to the Responders. In particular, we study the mechanism by which the Responders are selected from the population in order to elicit the best deal.

The NM-MUG can be used to study social settings in which groups of people wish to negotiate a deal. For example, deals between companies or between national legislative bodies are often discussed by committees representing each side, and as a result the selection of committee members with specific individual strategies can have a great influence on the final result (Hagan et al., 2001). Multiplayer versions of the Ultimatum Game are also played in the context of group buying (Jing and Xie, 2011). A specific example motivating our work are smart energy communities, such as the Amsterdam pilot sites *Schoonschip*¹ and *De Ceuve*², in which a number of households share a single point of coupling with the national energy grid. Behind this point

¹<http://schoonschipamsterdam.org/en/>

²<https://www.jouliette.net/>

of coupling, the households can exchange energy flexibilities (demand and supply) locally and thus more efficiently (Chakraborty et al., 2018). The summed remaining flexibility of each community could be used in negotiation with a different community, as a second layer of local or regional energy exchange (Lezama et al., 2018). This suggests a multiplayer bargaining game between two groups (the Proposing and Responding communities) which fits well within the general layout of the NM-MUG.

We simulate this scenario by means of a co-evolutionary process in which groups of Proposers and Responders are repeatedly selected from separate populations. The NM-MUG is used to compute the resulting fitness of individuals in each population, which then evolve following imitation dynamics and mutation. When selecting randomly composed groups of Proposers and Responders (that is, each individual has an equal probability of being selected for the group of Proposers or Responders), we find that the average offer of Proposers and acceptance thresholds of Responders co-evolves to an unfair state where Proposers get (almost) all the share. From this baseline, it is possible to test mechanisms for selecting the Responder groups, aiming to find arrangements that equalize the average gains of both populations. This model can thus lay out directions for future research in the areas social choice, group bargaining, competition, and the emergence of fairness in co-evolving communities.

Background and Related Work

The **Ultimatum Game (UG)** is a well-known interaction paradigm, widely used to evince the conflict between payoff maximization and fairness — and the puzzling human preference for the latter (Güth et al., 1982). As mentioned in the previous section, in this game two players interact in two distinct roles. One is called the *Proposer* and the other is denominated *Responder*. The game is composed by two sub-games, one played by each role. First, some amount of a given resource, e.g. money, is conditionally endowed to the Proposer; this agent must then suggest a division with the Responder. Secondly, the Responder will accept or reject the offer. The agents divide the money as it was proposed, if the Responder accepts. By rejecting, none of them will get anything. The strategy set of the Proposers comprises any possible division of the resource. The strategies of the Responders are acceptance or rejection, contingent on the offer made. Often, Responders' strategies are assumed to be probabilities of acceptance that are non-decreasing on the offer made. Frequently it is assumed that any Responder decision is codified in a threshold of acceptance: below this threshold offers are rejected (i.e., accepted with probability 0) and above the threshold offers are accepted with probability 1 (Page et al., 2000).

While the UG is a sequential game usually expressed in extensive-form, by having Responders deciding on their thresholds of acceptance in advance we can also formalize

this interaction as a normal-form game. In either case, the rational behaviour in the UG can be anticipated using traditional game-theoretical equilibrium analysis. Of special interest in this setting is the *sub-game perfect equilibrium* (Osborne et al., 2004), which can be inferred by applying the method of backward induction. The Responder, facing the decision of rejecting (earn 0) or accepting (earn some money, even if a really small quantity), would always prefer to accept. Secure about this certain acceptance, the Proposer will offer the minimum possible, maximizing her own share, thus yielding the equilibrium in which both the offer and the acceptance threshold are as close as possible to zero.

The UG is a 2-person game, however, many real-world situations require bargaining within (and between) groups of individuals. Here we focus on a multiplayer extension of the ultimatum game in which a group of N Proposers offers a division of resource to a group of M Responders (NM-MUG). A previous formalization of **Multiplayer Ultimatum Game (MUG)**, close to the one that we follow here, was proposed by Santos et al. (2015). In that work, a single Proposer makes an offer to a group of Responders. Individually, each Responder in the group states acceptance or rejection; the group of Responders as a whole accepts the offer provided that a minimum number of acceptances exist. A more recent study resorts to reinforcement learning (the Roth-Erev algorithm) to show that higher proposals are likely to emerge when stricter group decision rules (requiring more accepting Responders for group acceptance) are considered (Santos et al., 2016), also in the context of 1 Proposer versus N Responders. An alternative multiplayer (3-person) formulation of the UG was proposed by Takesue et al. (2017). Also, in a seminal work, Fehr and Schmidt (1999) explicitly considered the effect of competition between Proposers and Responders in a market game closely related with the UG. In this game, either 1) a group of sellers (Proposers) compete to sell one unit of a good to a buyer (Responder); or one Proposer suggests an offer that leads many Responders to compete against each other to accept it. In these market games, subjects tend to adopt unfairer strategies, differently to what happens with the 2-person UG and as predicted by the rationality self-interest model.

Nevertheless, both in the 2-person and the multiplayer ultimatum game, the predictions assuming perfect rationality were challenged by experimental and theoretical works (Fehr and Schmidt, 1999; de Jong and Tuyls, 2011; Santos et al., 2019). Instead of resorting to equilibrium notions of classical game theory to study the behavior of agents when interacting in a multi-Proposer multi-Responder ultimatum game, we adopt methods from population ecology, such as **evolutionary game theory (EGT)**. EGT has been used to analyze strategic interactions in several domains such as auctions (Phelps et al., 2004) or market dynamics (Bloembergen et al., 2015) (as an example). In a social context, EGT can describe individuals who revise their strategies

through social learning, being influenced by the behaviours and achievements of others (Sigmund, 2010). One of the most traditional tools to describe the dynamics of an evolutionary game model is the replicator equation (Taylor and Jonker, 1978). This equation, justified in a context of trait evolution in biology or cultural evolution across human societies, assumes that populations are infinite and evolution proceeds favouring strategies that offer a fitness higher than the average fitness of the population. However, it has also been argued that the replicator equation might not be an accurate model of human behaviour, due to its assumption of an infinite and well-mixed population, and that agent-based models might be more appropriate to study the social dynamics of fairness (Alexander, 2000).

Given these considerations, here we analyze the NM-MUG resorting to an agent-based model that similarly assumes that strategies performing better than average are selected over time. For that, we consider a pairwise comparison rule (Traulsen et al., 2006). As will be clarified below, we consider a finite population of agents. After playing several rounds, agents revise their strategy by observing a role-model agent, randomly picked from the population. Imitation (i.e., copying the strategy used by the role-model) occurs with a probability that grows with fitness difference: strategies performing better have a higher probability of being imitated. Under certain limits (large population size and low selection intensity) the replicator dynamics is recovered in this process (Traulsen et al., 2006).

NvM-Person Ultimatum Game

Let us start by describing the NvM-person (i.e., multi-Proposer, multi-Responder) Ultimatum Game, the interaction paradigm used throughout this paper. In any given NM-MUG interaction a group of Proposers makes an offer to a group of Responders. The offer made by the group results from a function of individual offers of Proposers in the group (e.g., the average); this offer is accepted if it is higher than a function of Responders' individual acceptance thresholds (e.g., if the offer is higher than the maximal threshold — guaranteeing that every Responder in the groups is satisfied — or if, again, the offer is higher than the Responders' average threshold). In case of acceptance, each Proposer receives the share she did not offer, which stresses the social dilemma in the Proposers' group: individually, each one has interest in offering the minimum possible but, in order to prevent a rejection, it is beneficial for all to have the largest possible collective offer.

In the context of smart energy communities, each group can be seen as a possible community, while the selection mechanism determines the attitude of its members. The proposals and response thresholds can be thought of as some combination of kWh and price, or the difference to the market price, illustrating the potential gain from the interaction.

Formally, we model the NM-MUG by two populations




Figure 1: The NvM-person Ultimatum Game (NM-MUG). Groups of N Proposers and M Responders are drawn from a population of Proposers Z_P and Responders Z_R , respectively. The groups' joint proposal and threshold for acceptance determine the success of the interaction.

Z_P and Z_R , representing the Proposers and Responders, respectively (see Figure 1). Each individual i in the population of Proposers is defined by her personal proposal value $p_i \in [0, 1]$, for $i \in Z_P$. Similarly, Responders are defined by their individual threshold of acceptance $q_j \in [0, 1]$, for $j \in Z_R$. At each iteration, a group of Proposers $N \subseteq Z_P$ and Responders $M \subseteq Z_R$ is selected, following predefined rules³. These groups induce a joint proposal $\bar{p} = \mathcal{P}(N)$ and joint Responder threshold $\bar{q} = \mathcal{Q}(M)$. In a simple scenario (such as the one we will consider below) \mathcal{P} and \mathcal{Q} are the average function, i.e. $\bar{p} = |N|^{-1} \sum_{i \in N} p_i$ and $\bar{q} = |M|^{-1} \sum_{j \in M} q_j$. The proposal is accepted iff $\bar{p} \geq \bar{q}$. The question is now: how to select the groups of Proposers and Responders from each population?

Base Scenario

In the base scenario, the Proposers, forming a group of fixed size N , are selected randomly from Z_P . The joint proposal offered by the group is taken to be the average proposal of individuals in the group, $\bar{p} = |N|^{-1} \sum_{j \in N} p_j$. The Responder group M is composed of those Responders that are willing to accept \bar{p} , such that $j \in M : q_j \leq \bar{p}$.

In this case, the Responders in M will have a payoff

$$U_i^R = \min(\bar{p}, \bar{p} \frac{N}{M}), \quad (1)$$

whereas all Responders outside M earn 0. The min operator signifies that the Responders cannot jointly receive more than what the Proposers offer, nor can one individual consume more than a unit share. At the same time, Proposer i taking part in the collective proposal by group N , offered to the group of Responders M , will have a payoff of

$$U_i^P = \min(1 - p_i, (1 - p_i) \frac{M}{N}), \quad (2)$$

³To simplify notation, we use N and M interchangeably as the group size of Proposers and Responders, respectively, as well as groups of selected Proposers and Responders. When an explicit distinction is necessary we use $|N|$ and $|M|$ to denote group sizes.

where p_i is the proposal by individual i . Again, the min operator ensures that Proposers cannot jointly offer more than the Responders accept; unit offers that are not accepted are lost in the context of the deal. This loosely reflects a typical scenario in which flexibilities are exchanged between smart energy communities (Lezama et al., 2018), where each individual household has a maximum amount of flexibility it can offer, and the total sum of flexibilities exchanged between the communities should balance out in the deal.

Responder Competition Scenario

As we detail below (Section “Experiments and Results”) allowing any individual $i \in Z_P$ to take part in the group of Responders (those that will accept or reject the offer and profit from it) has the pernicious effect of inducing a long-term reduction in the average values of q adopted in the Responders’ population which, in turn, incentivises the Proposers to lower their p and enact highly unfair offers. Many institutional arrangements affecting the process of Responder selection can be tested, departing from the base scenario presented above. For now, we discuss the role of Responder competition based on a declared threshold of acceptance — partly inspired by Fehr and Schmidt (1999). While Proposers are still randomly selected, we sort the Responders’ declared thresholds of acceptance, partitioning the Responders based on this ordering, and select for the group the individuals declaring the thresholds ranked from the m^{th} to the $(m + M - 1)^{th}$ ascending position. As an example of extreme cases, $m = 0$ and $M = 10$ means that the 10 lowest acceptance thresholds are selected and, in a population of 100 Responders, $m = 90$ and $M = 10$ means that the 10 highest acceptance thresholds are selected.

In this case, assuming that $|N|$ and $|M|$ are fixed externally, proposals are accepted only whenever $\bar{p} \geq \bar{q}$, where $\bar{q} = |M|^{-1} \sum_{j \in M} q_j$, and M is formed by the demands q_i which, after being sorted in an ascending order, stand in the positions ranging from the m^{th} to the $(m + M - 1)^{th}$ position. We study the evolutionary trajectories of strategy adoption when different rules for the selection of Responders are introduced (i.e., different m and M).

Evolutionary Dynamics

In order to study the evolutionary dynamics associated with each Responder selection rule (m and M), we implement an agent-based model in which individuals resort to social learning to adapt their behavior over time (Algorithm 1). Initially, values of p and q characterizing each agent are sampled from a uniform distribution. For a large number of generations, individuals will adapt their values of p and q . In each generation, $|Z_P| + |Z_R|$ individuals are sampled with replacement, following a uniform probability; with a probability μ the selected individual will randomly explore the strategy space, adopting a random value of p (if Proposer) or q (if Responder). This is akin to a mutation in genetic

Algorithm 1: Pseudo-code of the main cycle of our simulations. Algorithm 2 sketches how $fitness(\cdot)$ is computed.

```

Initialize all  $p_i \in Z_P, q_i \in Z_R = X \sim \mathcal{U}(0, 1)$ 
for  $t \leftarrow 1$  to  $Gens$  do Main cycle of interaction and
strategy update:
  for  $j \leftarrow 1$  to  $Z_P + Z_R$  do Select agent to update:
    if  $X \sim \mathcal{U}(0, 1) < Z_P / (Z_P + Z_R)$  then Update
      Proposer strategy:
        /* Sample two agents from
           Proposer population */
         $A \leftarrow X \sim \mathcal{U}(1, Z_P)$  (agent to update)
         $B \leftarrow X \sim \mathcal{U}(1, Z_P)$  (model agent)
      else Update Responder strategy:
        /* Sample two agents from
           Responder population */
         $A \leftarrow X \sim \mathcal{U}(1, Z_R)$  (agent to update)
         $B \leftarrow X \sim \mathcal{U}(1, Z_R)$  (model agent)
      if  $X \sim \mathcal{U}(0, 1) < \mu$  then Mutation:
         $p_A \leftarrow X \sim \mathcal{U}(0, 1)$ 
      else Imitation:
         $f_A \leftarrow fitness(A)$ 
         $f_B \leftarrow fitness(B)$ 
         $prob \leftarrow 1 / (1 + e^{-\beta(f_B - f_A)})$ 
        if  $X \sim \mathcal{U}(0, 1) < prob$  then
           $p_A \leftarrow p_B + \text{imitation error} \sim \mathcal{U}(-\epsilon, \epsilon)$ 

```

evolution. With probability $1 - \mu$ the selected individual A will resort to imitation. In this case, a model agent B from the same population is selected. The fitness of both agents is calculated as the average payoff obtained in a large number of NM-MUG interactions (Algorithm 2).⁴ Imitation will occur with a probability that follows the Fermi function for pairwise comparison (Traulsen et al., 2006) such that

$$prob_{A \leftarrow B} = \frac{1}{(1 + e^{-\beta(f_B - f_A)})}$$

where f_A is the fitness of the imitator, f_B is the fitness of the model, and β is the so-called intensity of selection, controlling how dependent the imitation process is on agents’ fitness values. When imitation occurs, the value of p or q characterizing agent B will be adopted by agent A . The adopted strategies are subject to a small perturbation: we add a value between $-\epsilon$ and ϵ , sampled from a uniform probability distribution. We guarantee that $p_i, q_j \in [0, 1], \forall i, j$, truncating the adopted value if necessary.

⁴Note: while Proposers are drawn independently of their strategy, Responders are selected based on their q_i which is why we need to take this into consideration in Algorithm 2.

Algorithm 2: Sketch of fitness computation of individual A based on selection of M Proposers and N Responders.

```

Function fitness( $A$ )
  accumulatedFitness = 0;
  for  $i \leftarrow 1$  to Samples do
    if  $A \in Z_P$  then Select Proposers including  $A$ :
       $\lfloor$  Sample  $|N| - 1$  other Proposers
    else
       $\lfloor$  Sample  $|N|$  Proposers
    Select group of Responders  $M$  (for instance,
      ordering their  $q$  values, ascending, and picking the
      agents having the thresholds in the range  $m^{th}$  to
       $(m + |M| - 1)^{th}$ 
       $\bar{p} = \sum_{j \in N} p_j / |N|$ 
       $\bar{q} = \sum_{k \in M} q_k / |M|$ 
      if  $\bar{p} \geq \bar{q}$  then Proposal accepted:
        if  $A \in Z_R \wedge A \in M$  then Compute Responder
          payoff:
             $\lfloor$  fitness  $\leftarrow U_A^R$  (using Equation 1)
          else Compute Proposer payoff:
             $\lfloor$  fitness  $\leftarrow U_A^P$  (using Equation 2)
          accumulatedFitness  $\ +=$  fitness
     $\lfloor$  return accumulatedFitness / Samples

```

During the simulations, we record 1) the average strategy used in the population of Proposers and Responders, 2) the average acceptance rate of proposals, 3) the average fitness of Proposers and Responders and 4) the time-series of strategy adoption. We are particularly interested in understanding how strategy dynamics are impacted by different Responder selection rules (i.e., different values of m and M). We report these results next.

Experiments and Results

We simulate the NM-MUG as described previously (Algorithm 1) and present the results in the following. We average over 100 runs of 20,000 generations each, and we use 100 samples for each fitness computation (Algorithm 2). We set $Z_P = Z_R = 100$, $\mu = 0.001$, $\epsilon = 0.01$, and $\beta = 10$.

We are interested in the fairness of accepted deals which, for simplicity, we here define as (expected) payoff equality within and between the populations. For our scenarios, this means that fairness *between populations* is achieved when

$$\bar{p} = \frac{|M|}{|M| + |N|} \quad (3)$$

such that proposers and responders on average receive the same utility (by Equations 1 and 2). Fairness *within populations* is similarly achieved when all individuals share the same expected utility, which in our scenario means that $p_i = p_j, \forall i, j \in Z_P$ and $q_i = q_j, \forall i, j \in Z_R$. While

more elaborate measures of fairness are possible (de Jong and Tuyls, 2011), we leave these for future work and instead focus on the role of Responder selection in this paper.

Base Scenario


After simulating the co-evolving dynamics of agents playing the NM-MUG, and adapting their p and q strategies accordingly, we realized that the base scenario, where all accepting individuals are selected, nurtures long-term unfair divisions between Proposers and Responders. We verified that the p and q evolve, on average — taken over the whole population(s), over 20,000 generations and over 100 runs — to values close to 0.01 and 0.1, respectively. This results in a large proportion of rejected deals and, in case a deal is accepted, a highly unfair (between populations) outcome. The same result is obtained when selecting randomly composed groups of Responders with a fixed size M .

Responder Competition

We proceeded to test how competition for taking part in the Responder group affects these dynamics. We measured the average strategy usage, acceptance rate and fitness given a range of values for the rank parameter m , yielding both extreme (very low or very high m) and moderate ($m \approx Z_R/2$) Responder groups.

As Figure 2a conveys, increasing m increases the average values of p and q adopted by individuals in the long-run. Notwithstanding, selecting strict Responders — that have the highest values of q — is pernicious by leading to low acceptance rates, as evidenced by Figure 2a (bottom panel). Selecting groups that are characterized by the lowest values of q (low m) is disadvantageous for the Responders population as, over time, Proposers learn to offer extremely low proposals. Selecting groups formed by the highest values of q (high m) is equally harmful: due to the high fraction of proposals being rejected, individuals are unable to obtain high values of fitness. The optimal selection rule selects those representatives with a value of q close to the population median (i.e., $m \approx Z_R/2$), as evidenced in Figure 2b. Note that this figure presents the *relative average fitness* as a ratio with respect to the case $m = 0$. For example, when $m = 45$ a Responder receives on average approximately 32 times more payoff compared to $m = 0$. Conversely, in this case a Proposer earns slightly above 50% of what she would receive for $m = 0$. Using this relative scale allows an easier comparison with the base (unfair) scenario $m = 0$. As $|N| = |M|$ in this case, a between populations fair outcome is achieved whenever $\bar{p} \approx 0.5$.


Changing m has a profound impact on the evolving dynamics of p and q , as shown in Figure 2c, where time-series corresponding to exemplifying runs for $m = 30$ (a), $m = 45$ (b), $m = 50$ (c), and $m = 90$ (d) are presented. We see a clear difference between situation (a), with high acceptance rates and low fairness, and (d), with almost arbitrary dynam-




(a) The average p , population- q (taken over the whole population) and group- q (taken over the selected groups of Responders). On the bottom we show the average acceptance rate, which is representative of the utilitarian social welfare in our scenario. Responder competition causes both values of p and q to increase when m increases (top). However, too large values of m result in an increasing number of rejected proposals (bottom).

Figure 2: Responder competition: Proposals are made by random groups of Proposers with size $N = 10$ and the group of Responders (with size $M = 10$) is formed by the Responders with the m^{th} to the $(m + M - 1)^{\text{th}}$ highest values of q .

ics due to the low acceptance rate and thus indiscriminate fitnesses within the population. Situation (b), with (on average) median ranked Responder groups yields high acceptance rates while simultaneously maintaining fair (between population) proposals. Interestingly, when the choice of qs to form the Responder groups is dictated by $m = 50$ (c), a cycling dynamic is often observed, representing periods of fairness and unfairness that repeatedly succeed over time. We extend the analyses of the average strategy used over time by presenting, in Figure 3, the full distribution of proposals (left) and acceptance thresholds (right) within those populations for $m = 45$ (top) and $m = 50$ (bottom). We find that the distribution of strategies is kept close to the mean, thus suggesting high within-population levels of fairness. In addition, we confirm that the cycles observed in Figure 2c for $m = 50$ do not result from the populations



(b) The fitness of Responders is maximized when intermediate groups of Responders (i.e., with the median values of q) are selected to form Responders' group. We represent the relative average fitness, as a ratio taken over the base scenario $m = 0$.



(c) Example of time series for scenarios marked in panels (a) and (b).

being divided in groups with different strategies that grow and shrink alternately — instead all agents adopt a similar strategy throughout.

Effect of Proposer Group Size

Finally, we investigate the effect of increasing the Proposers' group size, N . As hypothesized, increasing N yields a stricter social dilemma for the Proposers, akin to a public goods game: individuals will maintain a low value of p , expecting to maximize their share while hoping that others propose an offer high enough to guarantee acceptance by the Responders. As observed in Figure 4, this dilemma is more pressing in larger Proposer groups, as the average p adopted decreases with N . We also plot the between-population fair proposal value (black dashed line). By increasing the Proposer group size, the minimum proposal required to ensure a fair offer is relaxed: for the same Responders' group size

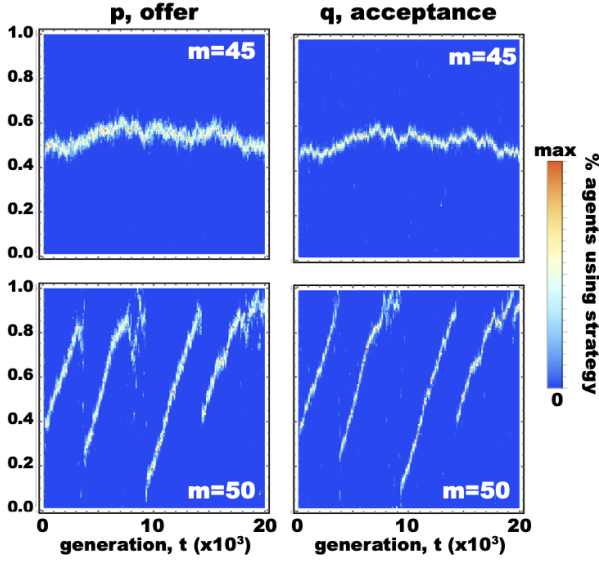


Figure 3: Distribution of strategies within the population for scenarios *b* (top) and *c* (bottom) highlighted in Figure 2. We find that agents adopt strategies that are close to the mean in each population, which implies high levels of within-population fairness.

(M), increasing N means that, individually, each Proposer is required to offer less, in order to maintain an even division between all Proposers and Responders involved in the transaction. In general, we find that fair proposals are easier to obtain when M is low and m is high. The average offer decreases whenever Proposers organize in larger groups.

Conclusion

In this paper we investigated the evolution of fairness within a new multiplayer version of the Ultimatum Game, the NM-MUG, in which a group of N Proposers offers a division of resources to a group of M Responders. Agent-based simulations showed that, in the absence of any Responder selection mechanism, the co-evolving populations of Proposers and Responders converge to a state in which proposals and acceptance thresholds are low, leading to unfair outcomes. This effect is more pronounced when the Proposers' groups are larger. We then investigated different Responder selection rules, contingent on their declared acceptance thresholds. We found that selecting extreme individuals is detrimental to the Responders' long-term payoff: selecting the lowest-demanding Responders incentives Proposers to do low offers, while selecting the highest-demanding Responders leads to many rejected offers. Moderate groups — i.e., selecting Responders with acceptance thresholds close to the population median — elicit the highest long-term gains for the Responders population as a whole, with high levels of fairness both between and within populations.

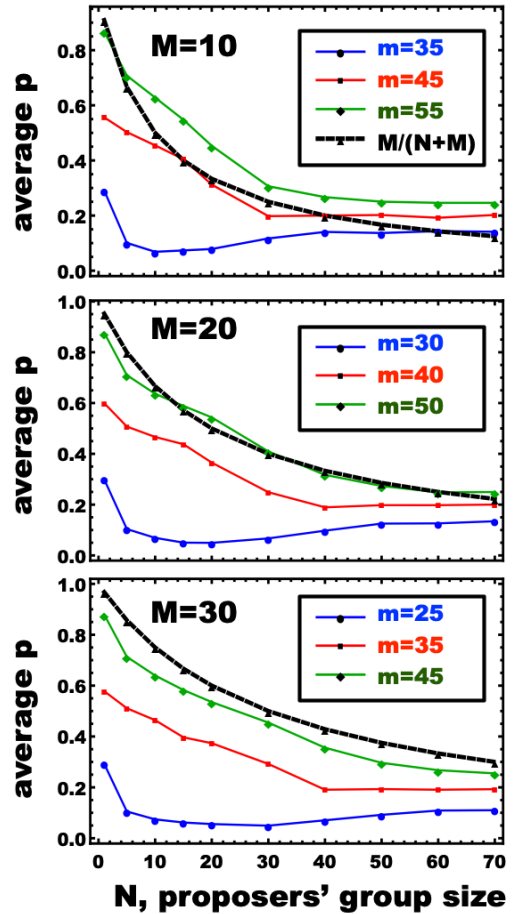


Figure 4: Effect of increasing Proposers' group size N on the average proposal level \bar{p} . We plot \bar{p} for Responders' group sizes $M = 10$ (top), $M = 20$ (middle), and $M = 30$ (bottom). Inside each panel, different Responders' selection rules (m) are plotted, moreover the black dashed line represents the between-populations fair proposal value (Equation 3), given the Proposers' and Responders' group sizes.

Although our model is based on several simplifying assumptions, due to its generality we believe that our results can nonetheless provide a practical message for institutional design (Bullock, 2016), both in the context of regional flexibility exchange between smart energy communities as well as in the broader context of group bargaining. The proposed model allows testing policies and emergent behaviors on the intersection between social choice theory, group bargaining, competition, and fairness elicitation.

We see many interesting avenues for further research based on our findings. A variety of selection rules can be envisioned and tested, for both the Responders as well as for Proposers. In addition, the utility functions used can be further tailored to specific real-world scenarios. Also, more elaborate measures of fairness can be investigated and could

even be incorporated into the utility function directly (as in e.g. de Jong and Tuyls 2011) or explicitly assuming traits such as empathy (Szolnoki et al., 2012), yielding potentially more complex and interesting dynamics. Finally, the effect of different selection rules can be analyzed in the context of spatially arranged individuals (Page et al., 2000; Szolnoki et al., 2012) or assuming iterated (multiplayer) ultimatum games (Ichinose and Sayama, 2014).

Acknowledgements

We are grateful to Michael Kaisers for helpful comments and feedback on earlier versions of this paper. This project has received funding in the framework of the joint programming initiative ERA-Net Smart Energy Systems' focus initiative Smart Grids Plus, with support from the European Union's Horizon 2020 research and innovation programme under grant agreement No 646039. F.P.S. acknowledges support from the James S. McDonnell Foundation.

References

- Alexander, J. M. (2000). Artificial justice. In *Proceedings of Artificial Life VII*, pages 513–522.
- Axelrod, R. and Hamilton, W. D. (1981). The evolution of cooperation. *Science*, 211(4489):1390–1396.
- Bloembergen, D., Hennes, D., McBurney, P., and Tuyls, K. (2015). Trading in markets with noisy information: An evolutionary analysis. *Connection Science*, 27(3):253–268.
- Bullock, S. (2016). Alife as a model discipline for policy-relevant simulation modelling: Might worse simulations fuel a better science-policy interface? In *Proc of the 2018 Conference on Artificial Life*, pages 28–29. MIT Press.
- Chakraborty, S., Baarslag, T., and Kaisers, M. (2018). Energy contract settlements through automated negotiation in residential cooperatives. In *2018 IEEE SmartGridComm*, pages 1–6.
- de Jong, S. and Tuyls, K. (2011). Human-inspired computational fairness. *Auton. Agents Multi Agent Syst*, 22(1):103–126.
- Fehr, E. and Schmidt, K. M. (1999). A theory of fairness, competition, and cooperation. *The Quarterly Journal of Economics*, 114(3):817–868.
- Gintis, H. (2000). *Game theory evolving: A problem-centered introduction to modeling strategic behavior*. Princeton University Press.
- Güth, W., Schmittberger, R., and Schwarze, B. (1982). An experimental analysis of ultimatum bargaining. *Journal of Economic Behavior & Organization*, 3(4):367–388.
- Hagan, J. D., Everts, P. P., Fukui, H., and Stempel, J. D. (2001). Foreign policy by coalition: deadlock, compromise, and anarchy. *International Studies Review*, 3(2):169–216.
- Hofmann, L.-M., Chakraborty, N., and Sycara, K. (2011). The evolution of cooperation in self-interested agent societies: A critical study. In *Proc. of 10th Int. Conf. on AAMAS 2011*, pages 685–692.
- Ichinose, G. and Sayama, H. (2014). Evolution of fairness in the not quite ultimatum game. *Scientific Reports*, 4:5104.
- Jing, X. and Xie, J. (2011). Group buying: A new mechanism for selling through social interactions. *Management Science*, 57(8):1354–1372.
- Lezama, F., Soares, J., Hernandez-Leal, P., Kaisers, M., Pinto, T., and do Vale, Z. M. A. (2018). Local energy markets: paving the path towards fully transactive energy systems. *IEEE Transactions on Power Systems*.
- Osborne, M. J. et al. (2004). *An introduction to game theory*, volume 3. Oxford university press New York.
- Page, K. M., Nowak, M. A., and Sigmund, K. (2000). The spatial ultimatum game. *Proc. Royal Soc. B*, 267(1458):2177–2182.
- Phelps, S., Parsons, S., and McBurney, P. (2004). An evolutionary game-theoretic comparison of two double-auction market designs. In *International Workshop on Agent-Mediated Electronic Commerce*, pages 101–114. Springer.
- Rand, D. G. and Nowak, M. A. (2013). Human cooperation. *Trends in Cognitive Sciences*, 17(8):413.
- Ranjbar-Sahraei, B., Bou Ammar, H., Bloembergen, D., Tuyls, K., and Weiss, G. (2014). Evolution of cooperation in arbitrary complex networks. In *Proc. of the 13th Int. Conf. on Autonomous Agents and Multi-Agent Systems*, pages 677–684.
- Santos, F. P., Pacheco, J. M., Paiva, A., and Santos, F. C. (2019). Evolution of collective fairness in hybrid populations of humans and agents. In *Proceedings of AAAI'19*. AAAI Press.
- Santos, F. P., Santos, F. C., Melo, F. S., Paiva, A., and Pacheco, J. M. (2016). Dynamics of fairness in groups of autonomous learning agents. In Osman, N. and Sierra, C., editors, *Autonomous Agents and Multiagent Systems*, pages 107–126.
- Santos, F. P., Santos, F. C., and Pacheco, J. M. (2018). Social norm complexity and past reputations in the evolution of cooperation. *Nature*, 555(7695):242.
- Santos, F. P., Santos, F. C., Paiva, A., and Pacheco, J. M. (2015). Evolutionary dynamics of group fairness. *Journal of Theoretical Biology*, 378:96–102.
- Sigmund, K. (2010). *The calculus of selfishness*, volume 6. Princeton University Press.
- Skyrms, B. (2004). *The stag hunt and the evolution of social structure*. Cambridge University Press.
- Szolnoki, A., Perc, M., and Szabó, G. (2012). Defense mechanisms of empathetic players in the spatial ultimatum game. *Physical Review Letters*, 109(7):078701.
- Takesue, H., Ozawa, A., and Morikawa, S. (2017). Evolution of favoritism and group fairness in a co-evolving three-person ultimatum game. *EPL (Europhysics Letters)*, 118(4):48002.
- Taylor, P. D. and Jonker, L. B. (1978). Evolutionary stable strategies and game dynamics. *Mathematical Biosciences*, 40(1-2):145–156.
- Traulsen, A., Nowak, M. A., and Pacheco, J. M. (2006). Stochastic dynamics of invasion and fixation. *Physical Review E*, 74(1):011909.