

Moderate Responder Committees Maximize Fairness in (N×M)-Person Ultimatum Games

Daan Bloembergen

Centrum Wiskunde & Informatica
Amsterdam, The Netherlands
d.bloembergen@cwi.nl

Fernando P. Santos

Department of Ecology and Evolutionary Biology
Princeton University, New Jersey, USA
fpsantos@princeton.edu

ABSTRACT

We introduce and study a multiplayer version of the classical Ultimatum Game in which a group of N Proposers jointly offers a division of resources to a group of 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 different smart energy communities, where the surplus of one community can be offered to meet the demand of a second community. We find that, in the absence of any mechanism, the co-evolving populations of Proposers and Responders converge to a state in which proposals and acceptance thresholds are low, as predicted by the rational choice theory. This is more evident if the Proposers' groups are larger (i.e., large N). Low proposals imply an unfair exchange that highly favors the Proposers. To circumvent this drawback, we test different committee selection rules which determine how Responders should be selected to form decision-making groups, contingent on their declared acceptance thresholds. We find that selecting the lowest-demanding Responders maintains unfairness. However, less trivially, selecting the highest-demanding individuals also fails to resolve this imbalance and yields a worse outcome for all due to a high fraction of rejected proposals. Selecting moderate Responders optimizes overall fitness. This result provides a practical message for institutional design and the model proposed allows testing policies and emergent behaviors on the intersection between social choice theory, committee selection and fairness elicitation.

1 INTRODUCTION

Many social dilemmas in society can be formulated and studied using game theoretic methods [8]. In particular, the question how cooperation can come about in a society of self-interested individuals has attracted considerable interest in the research community [1, 12, 15, 19, 20]. 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 [1, 25].

One example of such a game is the Ultimatum Game (UG) [10], in which one player, the Proposer, offers a certain split of a resource to a Responder, who decides to either accept or reject the offer.¹ If accepted, the players receive their share per the offer; if

rejected both players receive nothing. Here, we propose and study a Multiplayer version of the classical Ultimatum Game, in which a group of N Proposers jointly offers a division of resources to a group of M Responders. Henceforth we refer to this interaction as **(N×M)-person Multiplayer Ultimatum Game (NM-MUG)**. While a multiplayer version of UG was previously analyzed in the context of one Proposer and N Responders [21–23], 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 can be achieved, in which Proposers offer a substantial split to the Responders. In particular, we study the mechanism by which the committee of Responders is selected from the population in order to guarantee 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 [11]. Multiplayer versions of the Ultimatum Game are also played in the context of group buying [13]. A specific example motivating our work are smart energy communities, such as the Amsterdam pilot sites *Schoonschip*² and *De Ceuve*³, in which a number of households share a single point of coupling with the national energy grid. Behind this point of coupling, the households can exchange energy flexibilities (demand and supply) locally and thus more efficiently [5]. 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 [14]. This suggests a multiplayer bargaining between two groups (Proposing and Responding community) which fits well with the very general layout of NM-MUG.

We simulate this scenario by means of a co-evolutionary process in which committees of Proposers and Responders are repeatedly selected from separate populations. The NM-MUG is used to compute the resulting fitness of individuals in each populations, which then evolve following imitation dynamics and mutation. When selecting randomly composed committees of Proposers and Responders (that is, each individual has an equal probability of being selected for the group of Proposers or Responders in charge of negotiating), we

¹Although the Ultimatum Game is usually formulated as a sequential game, it can be cast as an identical normal form game in which the Responder decides on her response to any possible proposal in advance.

²<http://schoonschipamsterdam.org/en/>

³<https://jouliette.net/>

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 base scenario, it is possible to test mechanisms for selecting the Responder committees, aiming to find arrangements that equalize the average gains of both populations. This base model can thus lay out directions for future research in the areas social choice, committee selection and the emergence of fairness in co-evolving communities.

2 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 [10]. 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 [16]. As previously mentioned, while the UG is a sequential game usually expressed in extensive-form, by having Responders declaring their thresholds of acceptance we can also formalize this interactions as a normal-form game. In any 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* [17]. If one divides the UG in two stages it is possible to apply the method of backward induction to infer such an equilibrium. The last agent to play is the Responder: Facing the decision of rejecting (earn 0) or accepting (earn some money, even if a really small quantity), this agent would always prefer to accept. Secure about this deterministic acceptance, the Proposer will offer the minimum possible, maximizing his own share. Denoting by p the fraction of the resource offered by the Proposer, $p \in [0, 1]$, and by q the acceptance threshold of the Responder, $q \in [0, 1]$, acceptance will occur whenever $p \geq q$ and the *sub-game perfect equilibrium* of this game is defined by values of p and q slightly above 0.

While the UG is a 2-person game, there are several real-world situations that consist in 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 in [23]. In that work, a single Proposer makes an offer to a group of Responders. Individually, each Responder in the groups 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 [22], also in the context of 1 Proposer versus N Responders. An alternative multiplayer (3-person) formulation of the UG was proposed in [26]. Also, in a seminal work, Fehr and Schmidt explicitly considered the effect of competition between Proposers and Responders in a market game closely related with the UG [7]. 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 [6, 7, 9, 21]. 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 [18] or market dynamics [2] (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 [24]. One of the most traditional tools to describe the dynamics of an evolutionary game model is the replicator equation [27]. 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. The fact that replicator equation describes a process of social learning does not prevent it from being convenient in the understanding of individual learning [29]. A lot of effort has been devoted to bridging the gap between replicator dynamics and multiagent learning [3], especially after Börgers and Sarin showed that there is indeed an equivalence between replicator dynamics and a simple reinforcement learning algorithm (Cross learning) [4].

Inspired by EGT and the replicator dynamics, 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 [28]. As will be clarified below, we consider a population of many 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 (high population size and low selection intensity) the replicator dynamics is recovered in this process [28].

3 (N×M)-PERSON ULTIMATUM GAME

Let us start by describing the $(N \times M)$ -person (i.e., multi-Proposer, multi-Responder) Ultimatum Game, the interaction paradigm used




Figure 1: The (N×M)-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.

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.

Formally, we model the NM-MUG by two populations 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 rules⁴. 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?

3.1 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

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

⁴To 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.

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

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

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 [14], 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.

We are interested in the fairness of accepted deals, which we here define as payoff equality within and between the populations.⁵ For this base scenario this means that a **between communities** fair proposal is defined as $p^* = \frac{|M|}{|M|+|N|}$; a **within Proposers** fair proposal is defined as $p^* = p_i = p_j, \forall i, j$; and a **within Responders** fair proposal is obtained whenever $M = Z_R$. The results we discuss mainly stress between communities fairness, however we plan to explicitly consider, in future work, both the Proposers and Responders within population definitions of fairness.

3.2 Responder Competition Scenario

As we detail below (Experiments and Results, Section 4) allowing any individual $i \in Z_R$ to take part in the committee of Responders (those that will accept or reject the offer) has the pernicious effect of inducing a long-term reduction in the average values of q adopted in the Responders’ population which, in turn, incentives the Proposers to lower their p and enact highly unfair offers. Many institutional arrangements affecting the process of Responders committee 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 in [7]. 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 committee 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 that M is formed by the demands q_i that, after sorted in an ascending order, stand in the positions ranging from the m^{th} to the $(m + M - 1)^{th}$ positions. We study the evolutionary trajectories of strategy adoption when different rules for the selection of Responder committees are introduced (i.e., different m and M).

⁵More elaborate measures of fairness are possible [6], but left for future work.

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:
       $f_A \leftarrow fitness(A)$ 
       $f_B \leftarrow fitness(B)$ 
       $prob \leftarrow (1 + e^{-\beta(f_B - f_A)})^{-1}$ 
      if  $X \sim \mathcal{U}(0, 1) < prob$  then
         $p_A \leftarrow p_B + \text{imitation error} \sim \mathcal{U}(-\epsilon, \epsilon)$ 

```

3.3 Evolutionary Dynamics

In order to study the evolutionary dynamics associated with each Responder committee 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 evolution. With probability $1 - \mu$ the individual will resort to imitation. In this case, a model agent 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). Imitation will occur with a probability $(1 + e^{-\beta(f_B - f_A)})^{-1}$, 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. In case of imitation, the values of p or q characterizing agent B will be adopted by agent A . When imitation occurs, 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 strategies remain lower than 1 and greater than 0, truncating the adopted value if necessary.

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  $A$ :
      | Sample  $|N| - 1$  other Proposers
    else
      | 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:
        |  $fitness \leftarrow U_A^R$  (using Equation 1)
      else Compute Proposer payoff:
        |  $fitness \leftarrow U_A^P$  (using Equation 2)
      accumulatedFitness += fitness
  return accumulatedFitness / Samples

```

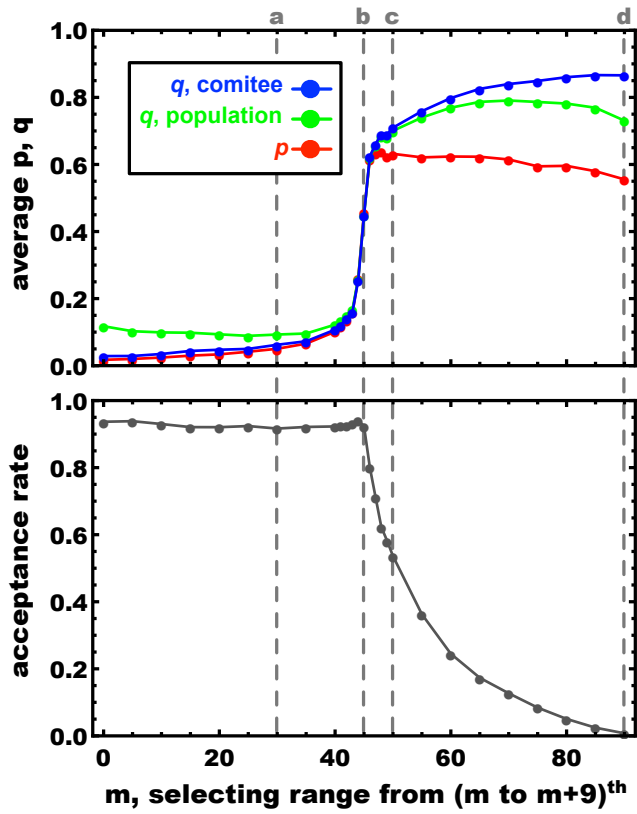
During the simulations, we record 1) the average strategy used in the population of Proposers and Responders, 2) the average acceptance rate of proposals, 3) the average fitness of Proposers and Responders and 4) the time-series of strategy adoption. We are particularly interested in understanding how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

4 EXPERIMENTS AND RESULTS

We simulate the NM-MUG as described previously and present the results in the following. We average over 100 runs of 20,000 generations each, and we use 100 samples for the fitness computation. We set $Z_P = Z_R = 100$, $\mu = 0.001$, $\epsilon = 0.01$, and $\beta = 10$. In future work we shall test the effect of varying β , although we expect that increasing this parameter will result in an overall decrease in fairness [19, 23].

After simulating the co-evolving dynamics of agents playing NM-MUG, and adapting their p and q strategies accordingly, we first realized that the **base scenario** 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.

We then proceeded to test how **Responder competition** for being included in the committee affects these dynamics. We measured the average strategy usage, acceptance rate and fitness given extreme (very low or very high m) and moderate ($m \approx Z_R/2$) Responder committees. As Figure 2a conveys, increasing m increases the average values of p and q adopted by individuals in the long-run. Notwithstanding, selecting extreme committees that have the highest values of q is (as evidenced by Figure 2a, bottom panel)

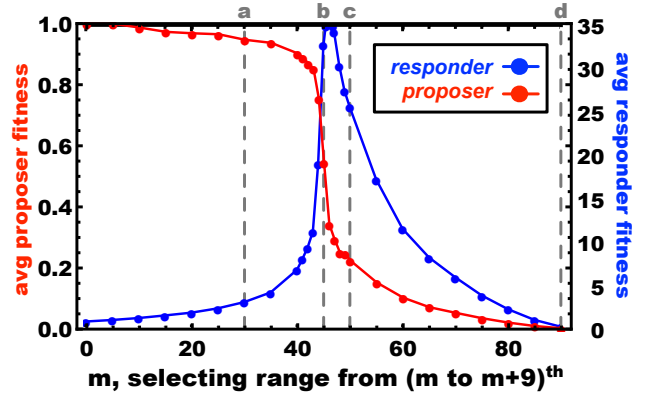


(a) The average p , population- q (taken over the whole population) and group- q (taken over the selected groups of Responders). On the bottom we show the average acceptance rate, which is representative of the utilitarian social welfare in our scenario. Responder competition causes both values of p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

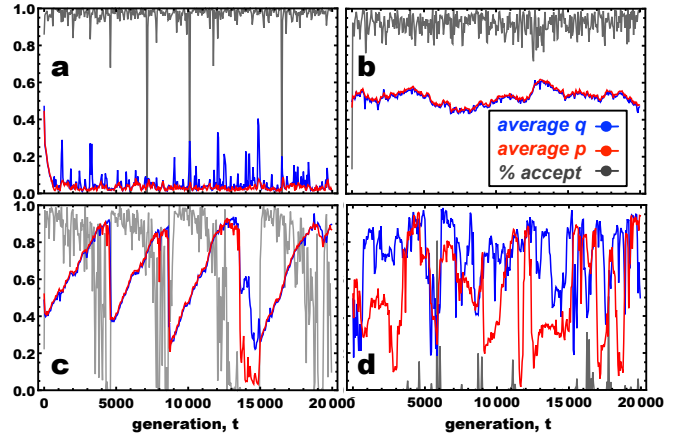
Figure 2: Responder competition: Proposals are made by random groups of Proposers with size $N = 10$ and the group of Responders is formed by the Responders with the m^{th} to the $(m + 9 - 1)^{\text{th}}$ highest values of q .

pernicious by leading to low acceptance rates. Selecting committees 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 committees 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 committee 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.

Changing m has a profound impact on the evolving dynamics of p and q , as represented in Figure 2c, where time-series corresponding to exemplifying runs for $m = 30$, $m = 45$, $m = 50$ and $m = 90$ are presented. Interestingly, whenever the choice of qs to form the committee is dictated by $m = 50$, a cycling dynamic is often observed, representing periods of fairness and unfairness that repeatedly succeed over time. In future work we shall delve



(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).

deeper into this question by further investigating the causes of such evolutionary dynamics.

Finally, we investigate the effect of increasing the Proposers' group size, N . As hypothesized before, increasing N yields a stricter social dilemma in the Proposer population, 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 3, this dilemma is more pressing in larger Proposer groups, as the average p adopted decreases with N .

5 CONCLUSION

In this paper we have proposed a new multiplayer version of the classical Ultimatum Game, the NM-MUG, in which a group of N Proposers jointly offers a division of resources to a group of M Responders. A preliminary study showed that, in the absence of

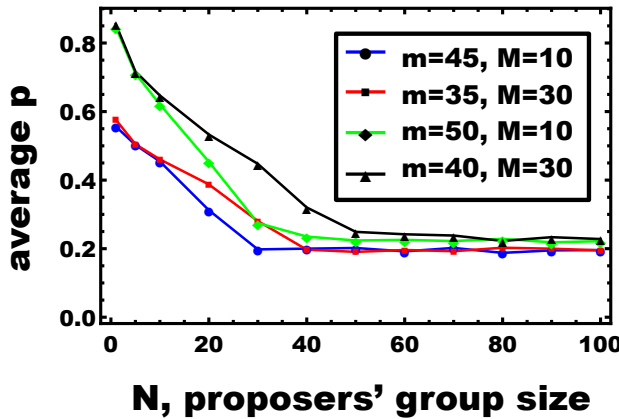


Figure 3: Effect of increasing Proposers' group size N on the average proposal level \bar{p} , for different Responders' selection committee rules (m). The average offer decreases – thus leading to a more unfair scenario – whenever Proposers organize in larger groups.

any group 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 committee selection rules, contingent on their declared acceptance thresholds. We found that selecting extreme committees is detrimental to the Responders' long-term payoff: selecting the lowest-demanding Responders incentivizes Proposers to offer low proposals whereas selecting the highest-demanding Responders leads to many offers being rejected. Moderate committees – 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. These first results provide a practical message for institutional design and the model proposed allows testing policies and emergent behaviors on the intersection between social choice theory, committee selection and fairness elicitation.

We see many interesting avenues for further research based on these first findings. Obviously, different committee selection rules can be envisioned and tested, for both the Responders as well as for Proposers. In addition, the utility functions used can be tailored to a specific real-world scenario such as the exchange of flexibilities between energy communities [14]. Furthermore, measures of fairness can be incorporated into the utility function directly (as in e.g. [6]), yielding potentially more complex and interesting dynamics.

ACKNOWLEDGMENTS

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

- [1] Robert Axelrod and William D. Hamilton. 1981. The evolution of cooperation. *Science* 211, 4489 (1981), 1390–1396.
- [2] Daan Bloembergen, Daniel Hennes, Peter McBurney, and Karl Tuyls. 2015. Trading in markets with noisy information: An evolutionary analysis. *Connection Science* 27, 3 (2015), 253–268.
- [3] Daan Bloembergen, Karl Tuyls, Daniel Hennes, and Michael Kaisers. 2015. Evolutionary dynamics of multi-agent learning: A survey. *Journal of Artificial Intelligence Research* 53 (2015), 659–697.
- [4] Tilman Börgers and Rajiv Sarin. 1997. Learning through reinforcement and replicator dynamics. *Journal of Economic Theory* 77, 1 (1997), 1–14.
- [5] Shantanu Chakraborty, Tim Baarslag, and Michael Kaisers. 2018. Energy Contract Settlements through Automated Negotiation in Residential Cooperatives. *arXiv preprint arXiv:1807.10978* (2018).
- [6] Steven de Jong and Karl Tuyls. 2011. Human-inspired computational fairness. *Autonomous Agents and Multi-Agent Systems* 22, 1 (2011), 103–126.
- [7] Ernst Fehr and Klaus M. Schmidt. 1999. A theory of fairness, competition, and cooperation. *The Quarterly Journal of Economics* 114, 3 (1999), 817–868.
- [8] Herbert Gintis. 2000. *Game theory evolving: A problem-centered introduction to modeling strategic behavior*. Princeton University Press.
- [9] Werner Güth and Martin G. Kocher. 2014. More than thirty years of ultimatum bargaining experiments: Motives, variations, and a survey of the recent literature. *Journal of Economic Behavior & Organization* 108 (2014), 396–409.
- [10] Werner Güth, Rolf Schmittberger, and Bernd Schwarze. 1982. An experimental analysis of ultimatum bargaining. *Journal of Economic Behavior & Organization* 3, 4 (1982), 367–388.
- [11] Joe D Hagan, Philip P Everts, Haruhiro Fukui, and John D Stempel. 2001. Foreign policy by coalition: deadlock, compromise, and anarchy. *International Studies Review* 3, 2 (2001), 169–216.
- [12] Lisa-Maria Hofmann, Nilanjan Chakraborty, and Katia Sycara. 2011. The Evolution of Cooperation in Self-Interested Agent Societies: A Critical Study. *Proc. of 10th Int. Conf. on AAMAS 2011* (2011), 685–692.
- [13] Xiaoqing Jing and Jinhong Xie. 2011. Group buying: A new mechanism for selling through social interactions. *Management Science* 57, 8 (2011), 1354–1372.
- [14] Fernando Lezama, Joao Soares, Pablo Hernandez-Leal, Michael Kaisers, Tiago Pinto, and Zita Maria Almeida do Vale. 2018. Local energy markets: paving the path towards fully Transactive energy systems. *IEEE Transactions on Power Systems* (2018).
- [15] Martin A. Nowak. 2006. Five rules for the evolution of cooperation. *Science* 314, 5805 (2006), 1560–1563.
- [16] Martin A Nowak, Karen M Page, and Karl Sigmund. 2000. Fairness versus reason in the ultimatum game. *Science* 289, 5485 (2000), 1773–1775.
- [17] Martin J Osborne et al. 2004. *An introduction to game theory*. Vol. 3. Oxford university press New York.
- [18] Steve Phelps, Simon Parsons, and Peter McBurney. 2004. An evolutionary game-theoretic comparison of two double-auction market designs. In *International Workshop on Agent-Mediated Electronic Commerce*. Springer, 101–114.
- [19] David G. Rand and Martin A. Nowak. 2013. Human cooperation. *Trends in Cognitive Sciences* 17, 8 (2013), 413.
- [20] Bijan Ranjbar-Sahraei, Haitham Bou Ammar, Daan Bloembergen, Karl Tuyls, and Gerhard Weiss. 2014. Evolution of cooperation in arbitrary complex networks. In *Proc. of Int. Conf. on AAMAS 2014*. International Foundation for Autonomous Agents and Multiagent Systems, 677–684.
- [21] Fernando P Santos, Jorge M Pacheco, Ana Paiva, and Francisco C Santos. 2019. Evolution of collective fairness in hybrid populations of humans and agents. In *Proceedings of AAAI'19*. AAAI Press.
- [22] Fernando P. Santos, Francisco C. Santos, Francisco S. Melo, Ana Paiva, and Jorge M. Pacheco. 2016. Dynamics of Fairness in Groups of Autonomous Learning Agents. In *Autonomous Agents and Multiagent Systems*, Nardine Osman and Carles Sierra (Eds.). Springer International Publishing, Cham, 107–126.
- [23] Fernando P. Santos, Francisco C. Santos, Ana Paiva, and Jorge M. Pacheco. 2015. Evolutionary dynamics of group fairness. *Journal of Theoretical Biology* 378 (2015), 96–102.
- [24] Karl Sigmund. 2010. *The calculus of selfishness*. Vol. 6. Princeton University Press.
- [25] Brian Skyrms. 2004. *The stag hunt and the evolution of social structure*. Cambridge University Press.
- [26] Hirofumi Takesue, Akira Ozawa, and So Morikawa. 2017. Evolution of favoritism and group fairness in a co-evolving three-person ultimatum game. *EPL (Europhysics Letters)* 118, 4 (2017), 48002.
- [27] Peter D Taylor and Leo B Jonker. 1978. Evolutionary stable strategies and game dynamics. *Mathematical biosciences* 40, 1-2 (1978), 145–156.
- [28] Arne Traulsen, Martin A Nowak, and Jorge M Pacheco. 2006. Stochastic dynamics of invasion and fixation. *Physical Review E* 74, 1 (2006), 011909.
- [29] Karl Tuyls and Simon Parsons. 2007. What evolutionary game theory tells us about multiagent learning. *Artificial Intelligence* 171, 7 (2007), 406–416.