Cost-Efficient Interventions for Promoting Fairness in the Ultimatum Game

Theodor Cimpeanu[1], Cedric Perret[2], and The Anh Han[1,⋆]

[1] School of Computing, Engineering and Digital Technologies, Teesside University, Middlesbrough, UK TS1 3BA

[2] College of Life and Environmental Sciences, University of Exeter, Exeter, UK EX4 4PY

⋆ Corresponding author: The Anh Han (T.Han@tees.ac.uk)

Abstract

Institutions and investors face the constant challenge of making accurate decisions and predictions regarding how best they should distribute their endowments. The problem of achieving an optimal outcome at a minimal cost has been extensively studied and resolved using several heuristics. However, these works usually failed to address how an external party can target different types of fair behaviour or do not take into account how limited information can shape this complex interplay. Here, we consider the Ultimatum game in a spatial setting and propose a hierarchy of interference mechanisms based on the amount of information available to an external decision-maker and desired standards of fairness. Our analysis reveals that monitoring the population at a macroscopic level requires more strict information gathering in order to obtain an optimal outcome and that local observations can mediate this requirement. Moreover, we identify the conditions which must be met for an individual to be eligible for investment in order to avoid unnecessary spending. We further explore the effects of varying mutation or behavioural exploration rates on the choice of investment strategy and total accumulated costs to the investor. Overall, our analysis provides new insights about efficient heuristics for cost-efficient promotion of fairness in societies. Finally, we discuss the differences between our findings and previous work done on cooperation dilemmas and present our suggestions for promoting fairness as an external decision-maker.

Keywords: Ultimatum game, interference, cost efficiency, decision making, evolutionary game theory, networks.

1 Introduction

The problem of how collective behaviour, such as cooperation, coordination, safety compliance and fairness among self-interested individuals, emerges in evolving, dynamical systems has fascinated researchers from many disciplines, ranging from Evolutionary Biology, Economics, Physics, Social Sciences and Computer Science (Airiau et al., 2014, Han, 2013, Han et al., 2020, Liu and Chen, 2020, Maynard-Smith, 1982, Nowak, 2006, Perc et al., 2017, Perc and Szolnoki, 2010, Sigmund, 2010, Tuyls and Parsons, 2007, West et al., 2007). Several mechanisms that are responsible for promoting the emergence of cooperation have been proposed, including direct and indirect reciprocity (Okada, 2020, Sigmund, 2010, Trivers, 1971), kin selection (Hamilton, 1964) and network reciprocity (Nowak, 2012, Ohtsuki et al., 2006, West et al., 2007). In these works, the evolution of desired collective behaviour is typically shaped by the combined actions of individuals within the systems.

On the other hand, external interference, where the advocating of certain desired collective behaviour is carried out by an external decision maker, who does not belong to the system, studies how this can be done in a cost-effective way (Chen et al., 2015, Cimpeanu et al., 2019, Han et al., 2018, Han and Tran-Thanh, 2018, Wang et al., 2019). These works aim to identify a broad class of interference strategies, or heuristics, that exploit available information such as global statistics (population behavioural composition), as well as local information such as local behaviour profile and diversity, and graph structures, for budget saving. This line of research is useful to provide insights into the design of self-organised and distributed Multi-Agent Systems (MAS), in order to ensure agents achieve a desired collective state. For instance, one might consider a hybrid system consisting of humans and intelligent machines, in which it is important to ensure a cooperative and trustful relationship amongst each other (Andras et al., 2018, Paiva et al., 2018, Santos et al., 2019). Another example is how international agencies such as the European Union and United Nations might advocate certain preferred political behaviours or resolve international conflicts, given a limited budget (e.g. in terms of cost and military resources) (Marton-Lafevre et al., 2007, Smidt, 2020).

The literature on external interference in evolving, dynamical systems (or populations) has so far focused on cooperation dilemmas, namely the Prisoner's Dilemma (PD) (Cimpeanu et al., 2019, Han et al., 2018, Han and Tran-Thanh, 2018) and the Public Goods Game (PGG) (Chen and Perc, 2014, Chen et al., 2015, Sasaki et al., 2012, Wang et al., 2019). In these games, the interactions are symmetric and the players' roles are equivalent. However, many real-world and

MAS interactions are asymmetric, where players may have different baseline characteristics and/or play different roles in the interactions (McAvoy and Hauert, 2015, Ogbo et al., 2021, Tuyls et al., 2018). Examples include conflict resolution (Selten, 1978, Smidt, 2020), technology adoption by firms (Ogbo et al., 2021), and multiparty resource allocations (Chevaleyre et al., 2005, Lerat et al., 2013), where participants might have different roles (e.g. proposers/dictators vs responders) or bargaining power in the decision making process. In this asymmetric setting, the external decision maker might need to take into account the difference among players' underlying characteristics, such as their roles in the interactions, in order to optimise the cost and the level of desired behaviour. In particular, we might ask, is it enough to target a subset of the roles to already achieve a sufficiently good outcome, since collecting information about all the roles might be (very) costly and time consuming?

This paper contributes to advancing the state-of-the-art by studying cost-efficient external interference in a spatial Ultimatum Game (UG), a popular bargaining game for investigating fair decision making in many disciplines, such as economics and AI/MAS research (Cimpeanu et al., 2021, de Jong and Tuyls, 2011, De Jong et al., 2008, de Melo et al., 2018, Fehr and Schmidt, 1999, Güth et al., 1982, Rauwolf and Bryson, 2018, Santos et al., 2019, Teixeira et al., 2021). While there is concern over which quantitative definition of fairness to use in several disciplines, such as algorithmic fairness in Artificial Intelligence and Machine Learning (Hutchinson and Mitchell, 2019, Mehrabi et al., 2021), the abstract framework of the UG presents an uncontroversial, unique mathematical criterion for measuring fairness. Fair individuals in this setting are those who choose to donate the larger proportion of their individual endowment. Moreover, the target variable and the model score are unrelated, and we test for each possible set of targets using this unique measure of fairness. Similar to previous works (Chen et al., 2015, Cimpeanu et al., 2019, Duong and Han, 2021a, Han et al., 2018, Han and Tran-Thanh, 2018, Wang et al., 2019), we measure cost-efficiency of an interference strategy by its total cost required over time. Furthermore, we will discuss explicitly the requirements for each strategy to be implementable, such as the availability of information about the population and local neighbourhoods, and budget constraints in each time step. In a standard UG, players have two different roles, proposer and receiver (or responder), with different bargaining powers (See Methods in Section 3 for a detailed description of the game). We consider the spatial version of the game (Page et al., 2000) where players are distributed on a network in order to examine how to exploit the roles' asymmetry in both global and local interference strategies (Han et al., 2018).

In general, a cost-efficient interference problem consists of solving a bi-objective optimisation problem (Han and Tran-Thanh, 2018, Wang et al., 2019), maximising the overall level of desired behaviours (in the long run) while ensuring the total cost spent being within budget and/or minimal. The key challenge is that, for evolving dynamical systems such as those in the above-mentioned examples, the system dynamics are shaped by various stochastic and random effects, such as those resulting from behavioural updates and mutation (behavioural exploration) (Rand et al., 2013, Traulsen et al., 2009). With behavioural updates, such as through social learning or reproduction (Nowak, 2006, Sigmund, 2010), undesired behaviours might resurface over time whenever interference was not sufficiently strong in the past. Through mutation, these behaviours might do so even when they were extinct. Hence, the external decision maker needs to take into consideration that they will have to repeatedly interfere in the system, in order to sustain the desired behaviour over time. Note however that, for simplicity, previous works have either omitted mutation (Han and Tran-Thanh, 2018, Wang et al., 2019), or assumed that it is infinitely small (for analytical treatment) (Duong and Han, 2021a, Han and Tran-Thanh, 2018). Mutation (behavioural exploration), where agents can freely experiment with new behaviours, is usually non-negligible in real populations and has been shown to play an important role in enabling cooperation in the context of social dilemmas (Antal et al., 2009, Duong and Han, 2019, 2021b, Han et al., 2012, Rand et al., 2013, Traulsen et al., 2009). Thus, the present work will also advance the state-of-the-art in this respect, where we will closely examine how different regimes of mutation, or agents' propensity for behavioural exploration, influence the manner in which external interference should be carried out. Indeed, our results show that when mutation is sufficiently high, in line with those observed in behavioural experiments, optimal interference strategies can be significantly different.

In the next section we describe the models and methods in detail, before presenting the results and a final discussion. We also include with this paper a Supplementary Information (**SI**) that includes additional results to support the robustness of the paper findings described in the main text.

2 Models and Methods

2.1 Ultimatum Game (UG)

Agents' interaction is modelled using the one-shot Ultimatum Game (UG) (Nowak et al., 2000, Page et al., 2000). In the UG, two players are offered a chance to win a certain sum of money, normalised to 1, which they must divide between each other. One player is elected proposer, and suggests how to split the sum, while the other, the receiver (responder) can accept or reject the deal. If the deal is rejected, neither player receives any part of the initial sum. As in (Nowak et al., 2000, Page et al., 2000), we assume that a player is equally likely to perform in one of the roles (proposer or receiver). A player's strategy is defined by a pair of parameters, p and q. When acting as proposer, the player offers the amount p, whereas in a receiver's role, the player rejects any offer smaller than q.

As we focus in this paper on the effect of having multiple roles on interference decision making, we consider a baseline UG model where proposers have two possible strategic offers, a low (L, with $p = l$) and a high one (fair) (H, with $p = h$), where $l < h \in [0, 1]$. On the other hand, receivers have two options, a low threshold (L, with $q = l$) and a high threshold (H, with $q = h$). Thus, overall, there are four possible strategies HH, HL, LH and LL (i.e. HL would denote proposing high and accepting any offers, etc.). Fairness is measured by calculating what percentage of the population is representative for either the HH or HL strategies (i.e., fair proposers), and this allows us to have a clear comparison with previous works—in terms of the level of population fairness achieved—that have studied the evolution of fairness in the UG, see e.g. (Nowak et al., 2000, Page et al., 2000, Rand et al., 2013). Unlike our work, they did not study the cost-efficiency of interference strategies for enhancing fairness.

Given evidence from several behavioural experiments (Güth et al., 1982, Rand et al., 2013), in which people (almost) never offered more than half of the sum in UG, we assume $h \leq 0.5$. Particularly, we set $h = 0.5$ and $l = 0.1$, as shown in (Page et al., 2000). In this scenario, the strategy LL is evolutionarily stable. We also confirm this result in our simulations, as shown in Figure S1 in Supplementary Information (SI) and we note that this result is true for several mutation rates.

The payoff for the four strategies HH, HL, LH and LL reads (for row player):

	HH	HL	LH	LL
HH	$\frac{1}{2}$	$\frac{1}{2}$	$\frac{1-h}{2}$	$\frac{1-h}{2}$
HL	$\frac{1}{2}$	$\frac{1}{2}$	$\frac{1-h+l}{2}$	$\frac{1-h+l}{2}$
LH	$\frac{h}{2}$	$\frac{1+h-l}{2}$	0	$\frac{1-l}{2}$
LL	$\frac{h}{2}$	$\frac{1+h-l}{2}$	$\frac{l}{2}$	$\frac{1}{2}$

2.2 Population structure and dynamics

We consider a population of agents or individuals on a square lattice of size $Z = L \times L$ with periodic boundary conditions— a widely adopted population structure in population dynamics and evolutionary games (Szabó and Fath, 2007). We focus our analysis on the efficiency of various interference strategies in spatial settings, adopting an agent-based model directly comparable with the setup of recent lab experiments on cooperation (Rand et al., 2014). We set $L = 100$ for all our experiments, resulting in a population size $Z = 10^4$. For the baseline results on well-mixed populations (complete graph), we chose a population size $Z = 100$.

Initially each agent is designated as one of the four strategies (i.e. HH, HL, LH, HH), with equal probability. At each time step or generation, each agent plays the UG with its (four) immediate neighbours. In the well-mixed baseline, each agent plays the UG with every other agent in the population. The score for each agent is the sum of the payoffs in these encounters. At the end of each generation an agent A with score f_A chooses to copy the strategy of a randomly selected neighbour agent B with score f_B with a probability given by the Fermi function (i.e. *stochastic update*) (Traulsen et al., 2006):

$$(1 + e^{(f_A - f_B)/K})^{-1},$$

where K denotes the amplitude of noise in the imitation process (Szabó and Fath, 2007). Varying K allows us to capture a wide range of update rules and levels of stochasticity, including those used by humans, as measured in lab experiments (Rand et al., 2013, Zisis et al., 2015). In line with previous works and lab experiments (Rand et al., 2013, Szabó and Fath, 2007, Zisis et al., 2015), we set $K = 0.1$ in our simulations. With a given probability μ, this process is replaced instead by a randomly occurring mutation. A mutation is equivalent to behavioural exploration, where the individual makes a stochastic decision switch to one of the

four available strategies. For each interference strategy, we study four different mutation rates, for $\mu \in \{10^{-4}, 10^{-3}, 10^{-2}, 2 * 10^{-1}\}$, as well as $\mu = 0$ for the deterministic update. We will explicitly state the values of mutation rates in all figures' captions. Note that we sometimes only include a certain subset of these results in the main text for clarity. For a comprehensive set of results, see SI.

Although our analysis below will focus on the stochastic update rule (in order to examine how stochasticity affects interference, as discussed above), we will also provide results for *deterministic update* to have a clear comparison with previous works (see e.g. (Han et al., 2018)). For the deterministic update, an agent's strategy is always changed to that of its highest scoring neighbour (Nowak and May, 1992, Szabó and Fath, 2007). This is a way of approximating the stochastic update rule where the stochastic effect is infinitely small, i.e. $K \to 0$.

We simulate this evolutionary process until a stationary state or a cyclic pattern is reached. Similarly to (Nowak and May, 1992), all the simulations in this work (described in next sections) converge quickly to such a state. For the sake of a clear and fair comparison, all simulations are run for 500 generations. Moreover, for each simulation, the results are averaged over the final 25 generations, in order to account for the fluctuations characteristic of these stable states. Furthermore, to improve accuracy, for each set of parameter values, the final results are obtained from averaging 30 independent realisations. When shown in figures, the error bars represent the standard error of the mean between replicates. In the case of the well-mixed baseline results, which we do not use for direct comparison of cost-efficiency, we run the simulations for 100 generations and average the results over the final 10 generations. This accounts for the faster convergence time characteristic of these networks, and the lower population size.

Note that in the special case of deterministic update (where we also do not consider mutations), simulations can stop early when the proportion of fair proposers reaches 100%. We note that when maximum fairness is not reached, investment can still be ongoing beyond 500 generations and thus, the total cost of interference is dependent on the chosen stopping point. However, our results show that the average investment at the 500 generation mark is never more than 0.2% of the average total investment, for all types of interference. Thus, this arbitrary number has a limited effect and should not affect these results qualitatively.

2.3 Cost-Efficient Interference in Networks

We aim to study how one can efficiently interfere in a structured population to achieve high levels of fairness while minimising the cost of interference. Naturally, the level of fairness is measured by the fraction of fair offers in the population (Rand et al., 2013), which is the total of HH and HL frequencies. An investment decision consists of a cost $\theta > 0$ to the external decision-making agent/investor, this value θ is added as surplus to the payoff of each suitable candidate. In order to determine cost-efficiency, we vary the individual investment amount θ for each proposed interference strategy, and we measure the total accumulated costs to the external investor. Thus, the most efficient interference schemes will be the ones with the lowest relative total cost.

We examine and compare different approaches of interference to induce fairness, based on ensuring fairness for either role or both, leading to different desirable behaviours to be targeted

(i) ensure all proposals are fair, thus investing in HH and HL (**Target: HH, HL**);

(ii) ensure only fair offers are accepted, thus investing in HH and LH (**Target: HH, LH**);

(iii) ensure both (i) and (ii), i.e. investing in HH only (**Target: HH**).

Moreover, in line with previous works on network interference (Chen et al., 2015, Cimpeanu et al., 2019, Han and Tran-Thanh, 2018), we will compare global interference strategies where investments are triggered based on network-wide information, and the local ones where investments are based on local neighbourhood information.

In the *population-based* approach, a decision to invest in desirable behaviours is based on the current composition of the population. We denote x_f the fraction of individuals in the population with a desirable behaviour, given a targeting approach at hand, i.e. (i), (ii) or (iii) as defined above. Namely, investment is made if x_f is less or at most equal to a threshold p_f (i.e. $x_f \leq p_f$), for $0 \leq p_f \leq 1$. They do not invest otherwise (i.e. $x_f > p_f$). The value p_f describes how rare the desirable behaviours should be to trigger external support. In the *neighbourhood-based* approach, a decision to invest is based on the fraction x_f calculated at local level. Investment happens if the proportion of neighbours of a focal individual with the desirable behaviours is less or at most equal to a threshold n_f (i.e. $x_f \leq n_f$), for $0 \leq n_f \leq 1$; otherwise, no investment is made.

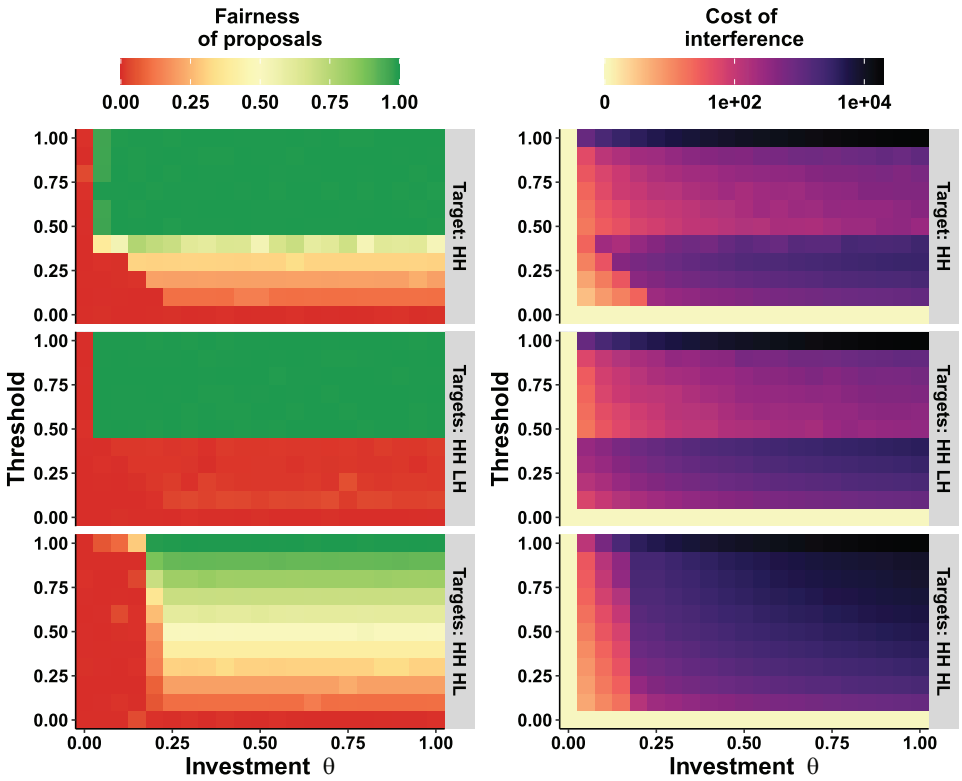

Figure 1. Baseline results after interference in a well-mixed scenario. Average fairness (left) and average cost of interference (right) as a function of the individual endowment θ and the threshold p_f (population-based, well-mixed network, $\mu = 0.01$, stochastic update). Each row represents a different targeting scheme. The cost of interference is shown on a logarithmic scale.

3 Results

When choosing to invest in a population of individuals in an effort to ensure some form of desirable outcome, an external decision maker must first consider several factors before any decision is made. Among these, we consider and aim to resolve the questions regarding what sort of behaviour they should invest in, how large the individual endowment must be, but also what an investor can do when information about the population or the environment is incomplete, or even unknown. As such, we consider that the simplest form of information gathering evaluates the overall population (in the form of some metrics measuring fairness on average), as opposed to fine-grained observations on individual neighbourhoods. Likewise, we consider that ensuring all proposals are fair (i.e. investing in HH or HL) is less demanding on an external decision-maker than ensuring that only fair offers are accepted (i.e. investing in HH and LH), which is, in turn, a simpler endeavour than for both the former and latter to be strictly enforced (choosing to invest in HH only). In this way, we can conceptualise a hierarchy of investment strategies, in terms of complexity, some of which may simply be impossible for an investor to follow, merely due to lack of information, funding, or a combination of the two.

We consider that there exists a *minimal level of fairness* which the external decision maker is aiming to enforce in regards to the population's behaviour (Han and Tran-Thanh, 2018), and we study the least expensive investment strategies for differing preferences of such an acceptable fairness.

3.1 Population-based results

Firstly, we explore the simplest class of investment strategies, using a macroscopic metric of the population, measuring average fairness in the whole system (population).

3.1.1 Baseline, Well-Mixed Populations

As the foundation of this analysis, we first introduce a baseline analysis of this interference on well-mixed (complete graph) populations, in Figure 1. We notice an increase in fairness for all three different targets, if the threshold for investment is sufficiently high, but there are marked differences in the cost of interference. Specifically, targeting both fair responses and proposals (HH), as well as only fair responses (HH LH), reduce the accumulated costs of interference for the external investor for a broader range of parameters than targeting only fair proposals (HH

HL). Furthermore, the threshold for investment is the deciding factor for ensuring high levels of fairness for all cases. This suggests that if certain levels of fair behaviour are maintained, then the population will converge to fairness without requiring further investment.

Based on the amount of information available to the external decision maker, we confirm that more information gathering leads to a more flexible investment approach. Respectively, the strictest approach (targeting HH only) leads to the highest levels of fairness with lowest accumulated costs, followed by ensuring fair responses, and, lastly, promoting fair proposals. Targeting both roles or only fair responses produce almost indistinguishable results if the chosen threshold is sufficiently high ($p_f \gtrapprox 40\%$), whereas only targeting fair proposers is very costly regardless of minimal fairness requirements. These results show that fair responders drive the dynamics of the system in the well-mixed scenario, and they should be targeted correspondingly by an external decision maker.

3.1.2 Structured Populations

We now consider that the population is structured and that individuals interact only with their neighbours. Figure 2 shows the results for different population-based interference scheme and clearly demonstrates the difference between the three targets for investment. We would like to point out the higher levels of fairness obtained using the HH targeting scheme, especially for a lower threshold p_f. We also notice an increase in the threshold for investment p_f in order to achieve similar levels of fairness. When it comes to the accumulated cost of interference, we see that HH is the most cost-effective solution, due to the previously perceived lower threshold required to maintain fairness.

Figure 3 further exemplifies the finding that targeting HH is the optimal scheme for population-based interference. Each row (portraying the different targeting schemes), drifts further away from the cost-optimal bottom left. As the threshold increases, so does the total cost, so the regions of high fairness for a lower threshold observed in Figure 2 coincide with the maximal savings (while still achieving desired levels of fairness).

Table 1 shows the most cost-efficient schemes for ensuring specific standards of fairness when only a population-based approach is possible, under differing rates of mutation (μ). We observe a definitive bias towards the most complex investment scheme (i.e. targeting HH players), which reiterates our previous observation. We note that, in order to maintain a desired level

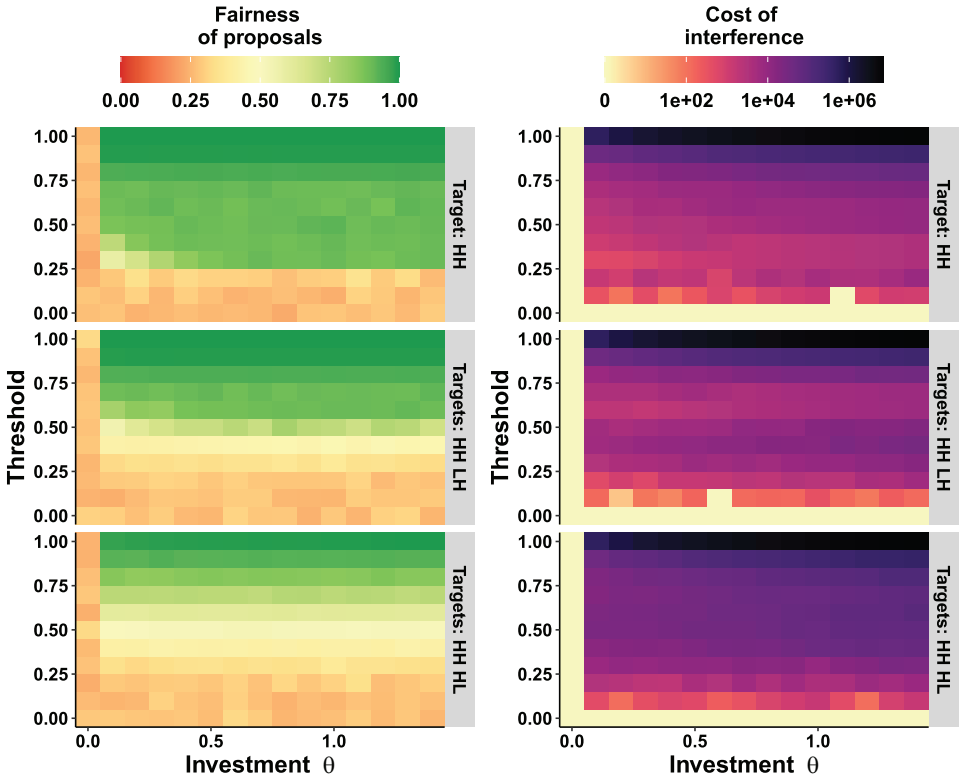

Figure 2. Average fairness (left) and average cost of interference (right) as a function of the individual endowment θ and the threshold p_f (population-based, $\mu = 0.01$, stochastic update). Each row represents a different targeting scheme. The cost of interference is shown on a logarithmic scale.

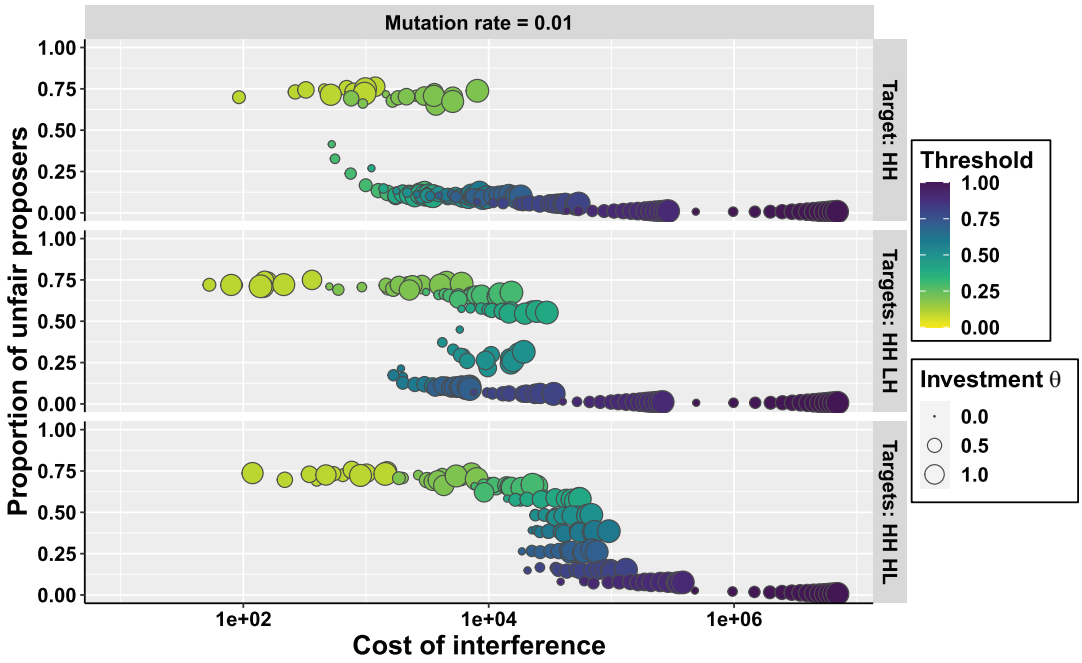

Figure 3. Proportion of unfair proposers as a function of average cost of interference for different targeting scheme (population-based, $\mu = 0.01$, stochastic update). The size and colour of the circles correspond to investment amount and threshold of investment, respectively. We note that the most desirable outcomes are closest to the origin.

Table 1. Most cost-efficient scheme to reach a minimum fairness of proposals for different mutation rates (population-based, stochastic update). There exist no schemes which satisfy the higher minimum fairness requirements in the case of very high mutation rate, written as '–' in the table.

Mutation rate	Minimum fairness	Target	Threshold	θ	Cost (mean $\pm$ 1.96 se)
10^{-4}	75%	HH	0.3	0.1	530 ± 5
10^{-4}	90%	HH	0.3	0.1	530 ± 5
10^{-4}	99%	HH	0.3	0.4	999 ± 7.6
10^{-2}	75%	HH	0.3	0.3	750 ± 5.4
10^{-2}	90%	HH	0.3	0.7	1747 ± 11.2
10^{-2}	99%	HH	1	0.1	487514 ± 93.6
0.2	75%	HH	0.6	0.2	358089 ± 650
0.2	90%	–	–	–	–
0.2	99%	–	–	–	–

of fairness, an external decision maker must increase the threshold at which they resume their investment, but also the individual endowment (θ). It becomes increasingly difficult to maintain standards of fairness when the population is exposed to high degrees of behavioural exploration and this naturally attracts an increase in total cost for the investor. We report similar figures for other values of μ in Figures S2, S3, S4, in SI.

Moreover, we observe an increase in fairness for all schemes of interference, across most values of individual endowment θ, which bodes well when the external decision maker possesses limited knowledge. If reducing cost is not the main objective, fairness can be maintained using any targeting scheme (i.e. any relevant observations made about the population), by increasing the minimum threshold p_f.

When the external decision maker is limited to the macroscopic metrics associated with population-based interference, interference is characterised by its strictness. To elaborate, information gathering should be the main goal for the investor, as ensuring that proposals and responses are simultaneously fair (i.e. targeting HH) is the optimal outcome. In this way, the minimum threshold can be kept low, reducing the accumulated cost. These findings are robust when compared to well-mixed populations, although it is easier for an investor to maintain fairness in the case of structured populations, when targeting fair proposers is the only option for investment.

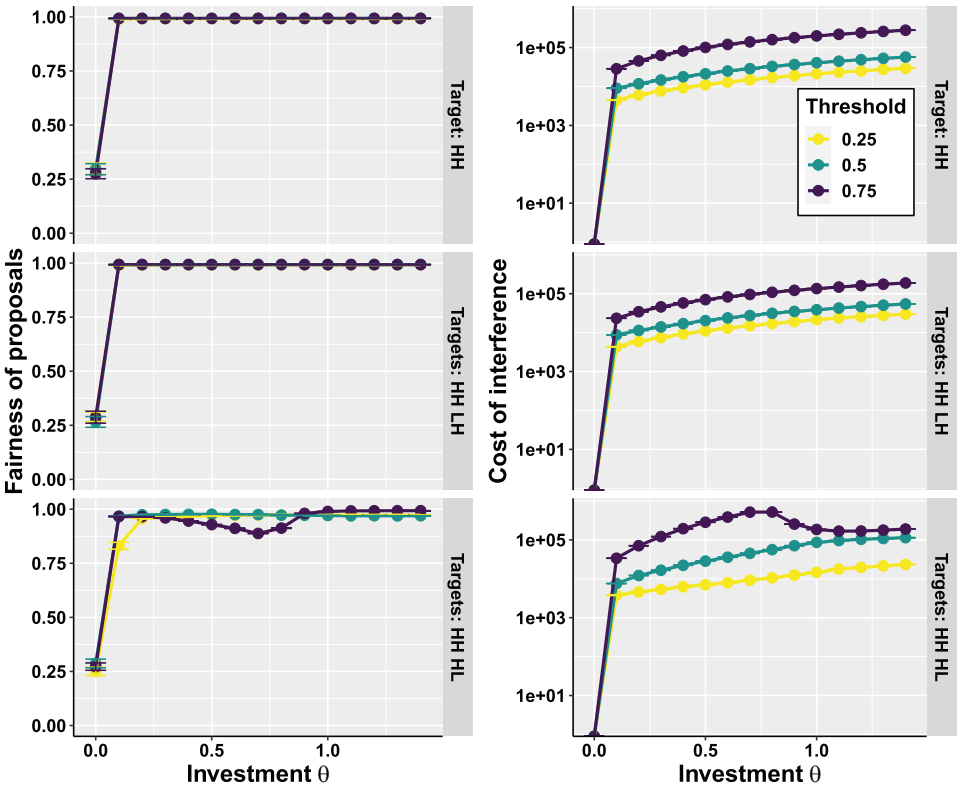

Figure 4. Average fairness (left) and average cost of interference (right) as a function of θ and threshold n_f (neighbourhood-based, $\mu = 0.01$, stochastic update). Each row represents a different targeting scheme. The cost of interference is on a logarithmic scale.

3.2 Neighbourhood-based interference

Previous works on the PD have shown that the greatest gains in cooperation (while maintaining a minimal investment cost) require very detailed observations of individual neighbourhoods, coupled with overly strict investment schemes (Cimpeanu et al., 2019, Han et al., 2018, Han and Tran-Thanh, 2018). In order to decipher whether or not these findings hold for the spatial ultimatum game, we study the outcome when an investor can perceive fairness at the local level.

Figure 4 reports the relationship between gains in fairness and increases in cost for an external investor, with diverse targets for receiving investments. We observe that fairness is more easily achieved than in population based interference, with only a very low investment required to sustain a majority of fair proposals. Further investment increases the cost of interference, but only slightly. If different thresholds result in fairness, Figure 4 shows that a threshold of 25% is the most cost-efficient. Similarly to population-based interference, the external decision maker should invest only when a large proportion of unfair individuals are present to limit the cost of investment. Finally, there are no significant differences between targeting schemes.

Similarly to our findings using a population-based approach, we observe that the more prohibitive option, HH, is also the most cost-effective. On the other hand, high fairness can be achieved in all three cases for the same values of endowment. Ensuring that all proposals are fair (thus investing in HH and HL), can lead to an increase in cost of interference, and a decrease in fairness gains (relative to the other two interference strategies). While all investment schemes evidently succeed in promoting the evolution of fairness, only ensuring the equitable proposals is not as reliable as encouraging discerning responses to offers or both. We note that this effect can only be seen when the threshold for investment is very high (i.e. an investor only invests in neighbourhoods with three or more fair proposers). As discussed earlier, investing in neighbourhoods with at most one fair agent and not investing otherwise, solves this dilemma.

Markedly, it is not effective to invest in neighbourhoods with a high percentage of fair proposals. These results point to a key observation, that it is more important to invest in fair proposers when there are few of them in a specific neighbourhood. In this sense, the lonely fair individuals require aid in otherwise competitive, unjust entourages. This result can further be seen in Figure 4. By being very selective with which neighbourhoods the external investor chooses to invest in (i.e. only choosing very fair neighbourhoods), they inadvertently produce a much higher final cost to their own selves. An external decision-maker would then unwittingly keep investing in fair proposals ad infinitum because fairness is eventually reached in the ultimatum game, even when individual endowment is relatively low. It is clear, therefore,

that to reduce potential costs, only players in unfair groups should be eligible for investment. Therefore, the defining characteristic of neighbourhood-based interference is the low threshold for investment (25%).

Table 2. Most cost-efficient scheme to reach a minimum fairness of proposals for different mutation rates (neighbourhood-based, stochastic update). There exist no schemes which satisfy the higher minimum fairness requirements in the case of very high mutation rates, written as '–' in the table.

Mutation rate	Minimum fairness	Target	Threshold	θ	Cost (mean $\pm$ 1.96 se)
10^{-4}	75%	HH	0.25	0.1	1395 ± 36.9
10^{-4}	90%	HH	0.25	0.1	1395 ± 36.9
10^{-4}	99%	HH	0.25	0.1	1395 ± 36.9
10^{-2}	75%	HH HL	0.25	0.1	3794 ± 200.1
10^{-2}	90%	HH LH	0.25	0.1	4352 ± 56.2
10^{-2}	99%	HH LH	0.25	0.2	5957 ± 60.7
0.2	75%	HH	0.25	0.4	150777 ± 121.5
0.2	90%	–	–	–	–
0.2	99%	–	–	–	–

By varying minimal fairness requirements and rates of mutation, we can gain further insight into which investment strategies are the most robust and cost-effective. Table 2 highlights some surprising findings. We see that neighbourhood based interference can result in a higher total cost than the optimal population-based interference schemes (see Table 1). Previous work has shown that more specific and restrictive intervention schemes are more effective in the PD (Cimpeanu et al., 2019, Han et al., 2018), but by being able to target different roles in the Ultimatum game, these differences can be mitigated. Furthermore, mutation rate serves as an equaliser between the investment targets, and we observe that less specific schemes (HH & HL and HH & LH) are the most cost-efficient options. We note that the differences between results are small enough that different runs could yield any outcome in the case of high or intermediate mutation rates. The lack of significant variability among the distinct targeting schemes contrasts strongly with the findings on the PD (Cimpeanu et al., 2019, Han and Tran-Thanh, 2018). We report similar figures for other values of μ in Figures S6, S7, S8, in SI.

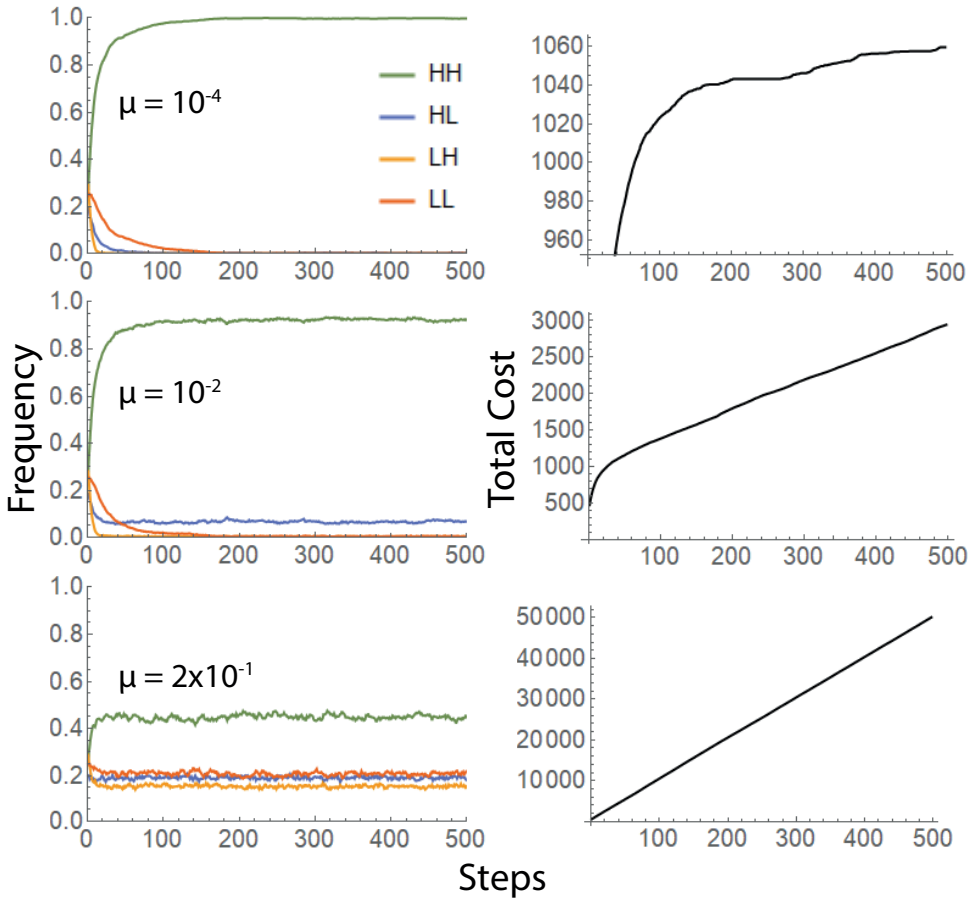

Figure 5. Typical runs showing the evolution of fairness and the associated total cost of interference for various mutation rates (top row $\mu = 10^{-4}$, middle row $\mu = 10^{-2}$, bottom row $\mu = 2 * 10^{-1}$; neighbourhood-based, stochastic update). Parameters: $n_f = 0.25,\ \theta = 0.1,\ Target = HH$. The choice of parameter values was motivated by selecting the optimal solutions in Table 2.

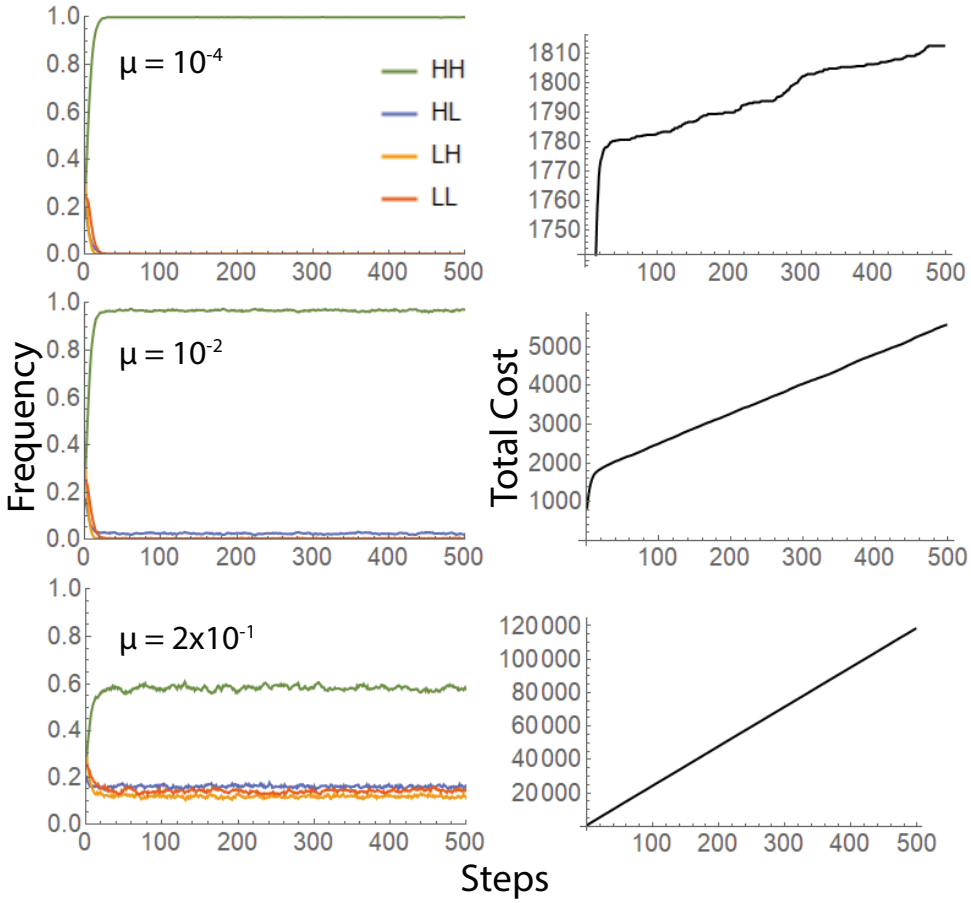

Figure 6. Typical runs showing the evolution of fairness and the associated total cost of interference for various mutation rates (top row $\mu = 10^{-4}$, middle row $\mu = 10^{-2}$, bottom row $\mu = 2 * 10^{-1}$; neighbourhood-based, stochastic update). Parameters: $n_f = 0.5,\ \theta = 0.1,\ Target = HH$. The choice of parameter values was motivated by selecting the optimal solutions in Table 2.

3.3 Evolution of strategies over time

We make use of the optimal parameter values identified in Tables 1 and 2 to explore the evolution of fairness over time for all the strategies in the population, as well as any associated accumulated costs. Through this analysis, we clarify some of the dynamics differentiating the different decisions for investment, as well as the effects of varying mutation rates upon the outcomes and the options available to investors.

The effects of mutation on the optimality of different interference schemes can be seen in Figure 5. As the mutation rate (μ) increases, the capacity of maintaining a threshold of fairness decreases (as also seen in Table 2). An external investor must increase their individual investment amount in order to meet these new demands set out by the increased mutation rates, and by doing so they can maintain fairness levels to a respectable standard.

To better highlight the sharp increases in the cost associated with the non-optimal threshold (i.e. when it is greater than 25%) for neighbourhood-based interference, we show such typical runs for varying mutation rates for the 50% threshold in Figure 6. When comparing Figures 5 and 6, we note the relative differences in total accumulated costs attributed to the choice of the threshold for investment n_f. We also note that increasing rates of behavioural exploration (mutation) amplifies this discrepancy.

We show how less specific interference strategies, which require less information gathering, can be effective in facilitating the evolution of fairness, when local monitoring is possible (Figure 8). Promoting fair proposals may often not be sufficient for low individual investment budgets (which are also the optimal solution) — in such cases fairness does not evolve. This occurs due to the inability of indiscriminate fair proposers to protect themselves against unfair proposers. Investing in fair proposers, in this case, artificially protects them against very competitive selection pressures.

Figure 7 showcases how different mutation rates call for different approaches to interference. As shown previously, optimal interference strategies vary according to the mutation rate. We point out the three different cases in which an investor might find themselves in. First, when few initial rounds of investment are enough for the system to converge and stabilise to a desired state. Second, an investor might be required to reinvest when the population tends to revert back to its initial condition. Lastly, constant investment is required to maintain a desired level of fairness, with the total cost skyrocketing accordingly. To some extent, a fair population can better deal with unfair invaders and this explains the need for a sufficiently high initial investment when mutation rates increase.

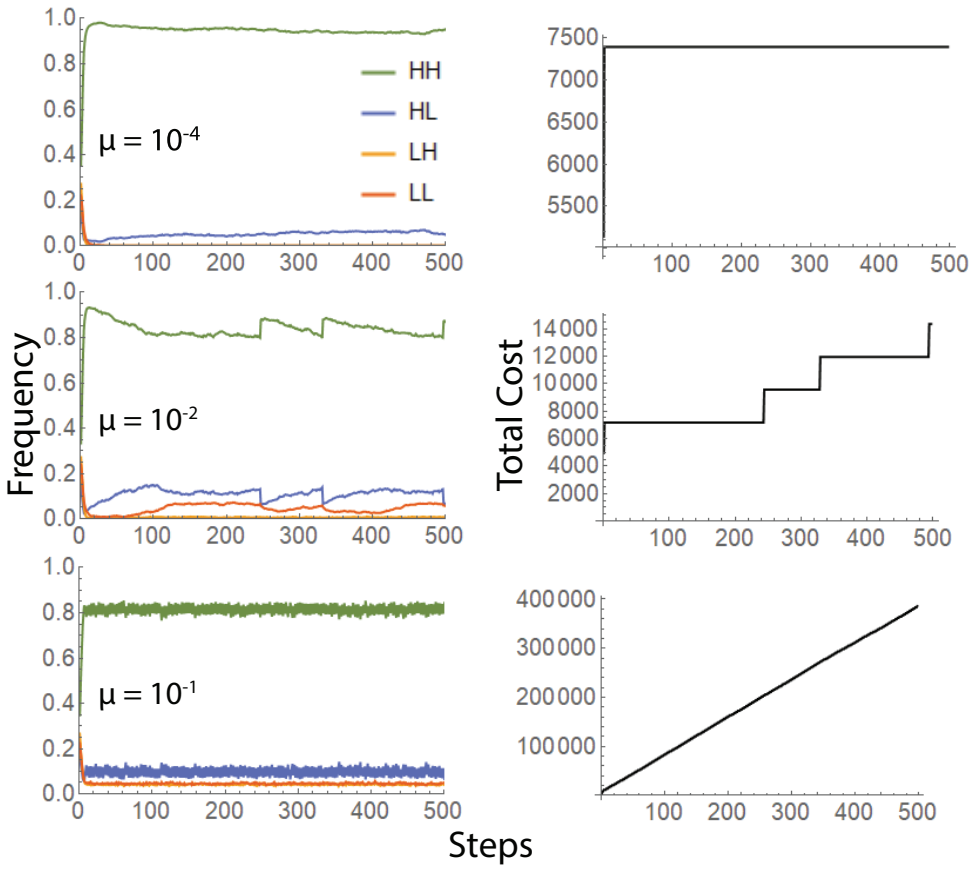

Figure 7. Typical runs showing the evolution of fairness and the associated total cost of interference for various mutation rates (top row $\mu = 10^{-4}$, middle row $\mu = 10^{-2}$, bottom row $\mu = 2 * 10^{-1}$; population-based, stochastic update). Higher mutation rates leads to an increasing need for interference over time. Parameters: $p_f = 0.8,\ \theta = 0.3,\ Target = HH$. The choice of parameter values was motivated by selecting the optimal solutions in Table 1.

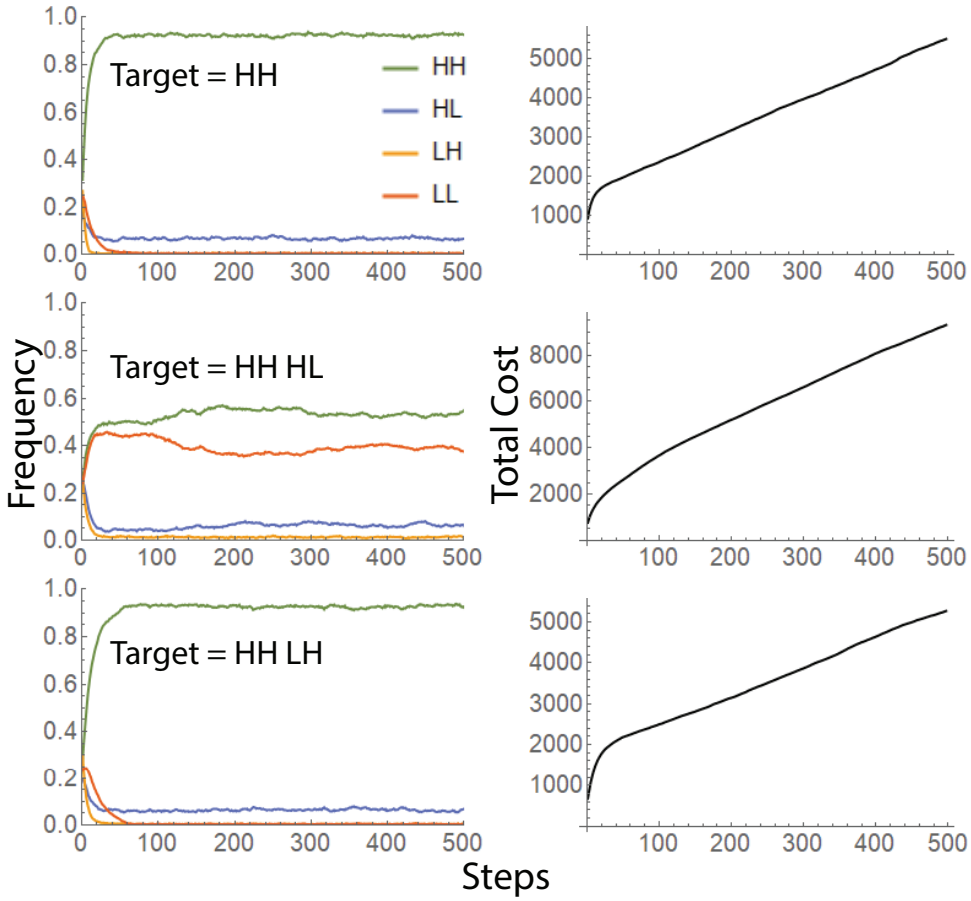

Figure 8. Typical runs showing the evolution of fairness and the associated total cost of interference for various targeting schemes (neighbourhood-based, stochastic update). Parameters: $n_f = 0.25,\ \theta = 0.2,\ \mu = 10^{-2}$. The choice of parameter values was motivated by selecting the optimal solutions in Table 2.

Finally, behavioural exploration motivates the manner or strength (in terms of individual endowment) of any initial efforts to moderate unfair behaviour. Figures 5, 6 and 7 show that the increase in cost is linear and ever-growing for high mutation-rates and gradually sharper at the beginning for lower mutation, eventually plateauing when the population is exposed to little or no behavioural exploration.

Table 3. Most cost-efficient population-based scheme (deterministic update) to reach a minimum fairness of proposals.

Minimum fairness	Target	Threshold	θ	Cost (mean $\pm$ 1.96 se)
75%	HH	0.5	0.5	1251 ± 10.8
90%	HH	0.6	0.9	2228 ± 22.6
99%	HH	0.9	1.1	5488 ± 22.9

Table 4. Most cost-efficient neighbourhood-based scheme (deterministic update) to reach a minimum fairness of proposals.

Minimum fairness	Target	Threshold	θ	Cost (mean $\pm$ 1.96 se)
75%	HH	0.25	0.8	2146 ± 56.3
90%	HH	0.25	0.8	2146 ± 56.3
99%	HH	0.25	1	2513 ± 16.4

3.4 Deterministic update

The results and findings reported so far were based on the stochastic update rule. We now take a step back and consider whether our findings would still hold for the deterministic rule (see again Methods section). It is not only for the sake of a direct comparison with a previous analysis reported in (Han et al., 2018), where cost-efficient interference was studied for the spatial PD in a deterministic setting (with no mutation). It would also allow us to examine if the findings above would remain robust for the deterministic update, a popular approximation for rare stochastic effect (or infinite intensity of selection) that is regularly used in the literature

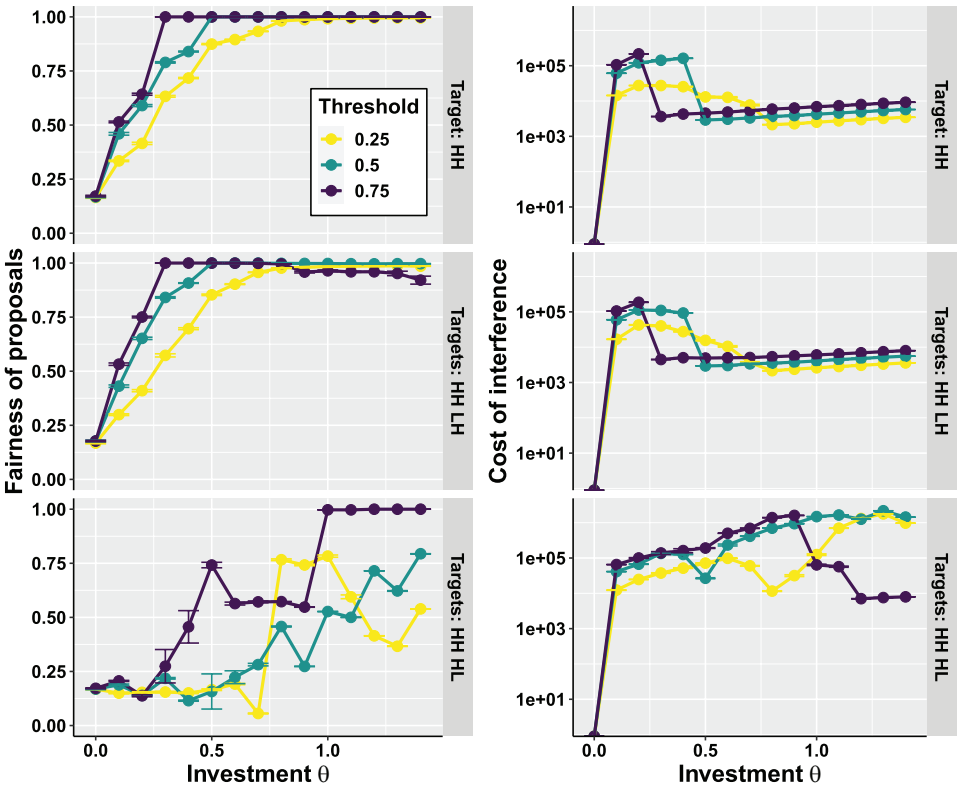

Figure 9. Average fairness (left) and average cost of interference (right) as a function of θ and threshold n_f (neighbourhood-based, deterministic update). Each row represents a different targeting scheme. The cost of interference is on a logarithmic scale for clarity.

(Szabó and Fáth, 2007). Tables 3 and 4 report results for the optimal interference strategies, in population-based (for a full report, see Figure S9, in SI) and neighbourhood-based schemes, respectively. We observe that for both schemes, targeting HH is always the best option. This is the same as the stochastic approach for population-based schemes but different from the neighbourhood-based ones. However, for the latter ones, the optimal threshold $n_f = 0.25$ remains the same as in the case of stochastic update, see also Figure 9. This is in stark contrast with the PD results where $n_f = 0.75$ was always the optimal choice.

4 Discussion

In summary, this paper has advanced the state of the art of the literature on external interference in dynamical systems, or populations of self-interested individuals, in two main respects: i) we have addressed an asymmetric interaction setting, in the form of the Ultimatum game, where players have different roles in the interaction. We have shown that it is crucial to consider the roles' asymmetry to provide cost-efficient investment strategies. This important analysis was not possible in previous works where symmetric games were studied (Chen and Perc, 2014, Chen et al., 2015, Cimpeanu et al., 2019, Han et al., 2018, Han and Tran-Thanh, 2018, Wang et al., 2019); ii) we have incorporated realistic levels of mutation or behavioural exploration in our analysis and have shown that they strongly affect the manner in which interference should be carried out. Previous works have always omitted mutation or assumed that it is infinitely small, thereby being unable to address this important issue for real-world populations and applications.

We have identified several key features that are required for a cost-effective interference scheme. On the one hand, population-based schemes are characterised by the need of extensive information gathering about both roles, as targeting HH always leads to the optimal strategy. On the other hand, neighbourhood-based schemes are characterised by their flexibility, where the optimal strategy always entails that investment is only made when there is at most one player with the desirable behaviour in the neighbourhood (i.e. no investment should be made when there is a half or larger fraction of such behaviour the neighbourhood). Our findings stand out in stark contrast with previous works on cooperation dilemmas, where both population and neighbourhood-based schemes require a highly strict investment approach.

The ultimatum game has been widely studied, whether with theoretical models (see Debove et al. 2016 for a review) or behavioural experiments (see Güth and Kocher 2014 for a review). The main motivation arises from the gap between theoretical predictions, in which rational indi-

viduals keep most of the endowment, and in which responders accept any positive proposition, however small it may be, and experiments, in which individuals propose 40% to 50% of their endowment (and often get punished if they propose less) (Güth et al., 1982). That said, previous works have investigated how fairness can evolve in models of the ultimatum game, wherein several mechanisms promoting the emergence of fairness have been identified. We note that we align our definition of fairness with these previous works, where generous proposers are deemed as fair, regardless of their behaviour when acting in the role of the responder.

Among others, Nowak et al. (2000) have studied the evolution of fairness in the Ultimatum game under indirect reciprocity, i.e. when players can observe others' interactions and have information about the reputations of others. We do not rely on reputation building mechanisms, as the role of this mechanism can be limited in large groups, where one-shot interactions between strangers are common. Page et al. (2000), Sinatra et al. (2009) have developed spatial models of the ultimatum game, where interactions happen only between neighbours. It has been shown that a spatial structure can promote the emergence of higher levels of fairness, but an equal split between proposers and responders is yet to be reached. The model developed in the present work has also considered a spatial model because (i) it captures an essential feature of many real-world networks of contacts (Barabasi, 2014), and (ii) it allows us to explore the effects of interference localised in particular neighbourhoods, which has been shown to be more cost efficient (Han et al., 2018). Furthermore, Rand et al. (2013) have shown that even if the population is well-mixed, fairness can result directly from the effect of randomness due to mutation and stochastic strategy updates. Intuitively, the uncertainty in the responders' choice forces the proposers to offer a high proposal so as to avoid rejection. We show that these stochastic factors also strongly influence the manner in which external interference may be performed while maintaining cost-efficiency. Moreover, the Ultimatum game has also been used to study how fairness can emerge in a hybrid population of human-agent interactions (de Melo et al., 2018, Santos et al., 2019). These works, however, do not consider external interference.

The problem of how to externally influence a system of multiple interacting agents to achieve a certain desired behaviour has been of significant interest in mechanism design, network theory and control theory literature. For example, how to maximise influence in networks has been studied in (Bloembergen et al., 2014, Riehl and Cao, 2016, Wilder et al., 2018). Moreover, Endriss et al. (2011) have investigated how to tax games so as incentivise certain behaviours at system equilibrium ; while Wooldridge (2012) has presented potential ways to manipulate games in order to achieve desired behaviours. These works, however, assume that

the decision-maker possesses full control of the agents within the systems. With our approach, the decision-makers have little or no direct control on the agents' behaviour, so they can rely only on rewarding schemes and their effects as ways of motivating the evolution of fairness. It is noteworthy that these works do not focus on the cost efficiency problem, whereas cost optimisation is one of our main goals.

Although the main focus of our discussion centres around structured populations in the four-neighbour lattice, our results show that similar trends can be observed in well-mixed populations, as well. When we consider population-based interference that can be enacted both on lattice and well-mixed populations, our observations are robust across network types, and we see that targeting both roles leads to the most cost-effective path towards fairness. In order to determine whether or not spatial heterogeneity plays a key role when multiple roles can be targeted by an external investor, future works will need to also consider more complex structures, such as scale-free or multiplex networks.

Our future work will examine other asymmetric games with multiple roles, such as the trust and anticipation games, where the bargaining nature is different from the ultimatum game (Gut, 2009, Han et al., 2021, Rauwolf and Bryson, 2018, Zisis et al., 2015), to see how this bargaining factor might affect the way interference needs to be made. We are also interested in how different network structures influence the interference strategies in asymmetric interactions, which has been studied for symmetric games (Cimpeanu et al., 2019).

Acknowledgements

Some preliminary results from this work was previously published as an extended abstract in AAMAS 2021 conference, pages 1840–1842, see ref. (Cimpeanu et al., 2021). T.C., C.P. and T.A.H. were supported by Future of Life Institute grant RFP2-154. T.A.H. is also supported by a Leverhulme Research Fellowship (RF-2020-603/9).

References

Airiau, S., Sen, S., and Villatoro, D. (2014). Emergence of conventions through social learning. *Autonomous Agents and Multi-Agent Systems*, 28(5):779–804.

Andras, P., Esterle, L., Guckert, M., Han, T. A., Lewis, P. R., Milanovic, K., Payne, T., Perret,

C., Pitt, J., Powers, S. T., Urquhart, N., and Wells, S. (2018). Trusting Intelligent Machines: Deepening Trust Within Socio-Technical Systems. *IEEE Technology and Society Magazine*, 37(4):76–83.

Antal, T., Traulsen, A., Ohtsuki, H., Tarnita, C. E., and Nowak, M. A. (2009). Mutation-selection equilibrium in games with multiple strategies. *Journal of theoretical biology*, 258(4):614–622.

Barabasi, A.-L. (2014). *Linked-how Everything is Connected to Everything Else and what it Means F*. Perseus Books Group.

Bloembergen, D., Sahraei, B. R., Bou-Ammar, H., Tuyls, K., and Weiss, G. (2014). Influencing social networks: An optimal control study. In *ECAI*, volume 14, pages 105–110.

Chen, X. and Perc, M. (2014). Optimal distribution of incentives for public cooperation in heterogeneous interaction environments. *Frontiers in behavioral neuroscience*, 8:248.

Chen, X., Sasaki, T., Brännström, Å., and Dieckmann, U. (2015). First carrot, then stick: how the adaptive hybridization of incentives promotes cooperation. *Journal of the royal society interface*, 12(102):20140935.

Chevaleyre, Y., Dunne, P. E., Endriss, U., Lang, J., Lemaitre, M., Maudet, N., Padget, J., Phelps, S., Rodrígues-Aguilar, J. A., and Sousa, P. (2005). Issues in multiagent resource allocation.

Cimpeanu, T., Han, T. A., and Santos, F. C. (2019). Exogenous rewards for promoting cooperation in scale-free networks. In *Artificial Life Conference Proceedings*, pages 316–323. MIT Press.

Cimpeanu, T., Perret, C., and Han, T. A. (2021). Promoting fair proposers, fair responders or both? cost-efficient interference in the spatial ultimatum game. In *In Proc. of the 20th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2021)*, pages 1480–1482.

de Jong, S. and Tuyls, K. (2011). Human-inspired computational fairness. *Autonomous Agents and Multi-Agent Systems*, 22(1):103–126.

De Jong, S., Uyttendaele, S., and Tuyls, K. (2008). Learning to reach agreement in a continuous ultimatum game. *Journal of Artificial Intelligence Research*, 33:551–574.

de Melo, C. M., Marsella, S., and Gratch, J. (2018). Social decisions and fairness change when people's interests are represented by autonomous agents. *Autonomous Agents and Multi-Agent Systems*, 32(1):163–187.

Debove, S., Baumard, N., and André, J. B. (2016). Models of the evolution of fairness in the ultimatum game: A review and classification. *Evolution and Human Behavior*, 37(3):245–254.

Duong, M. H. and Han, T. A. (2019). On equilibrium properties of the replicator–mutator equation in deterministic and random games. *Dynamic Games and Applications*, pages 1–23.

Duong, M. H. and Han, T. A. (2021a). Cost efficiency of institutional incentives in finite populations. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*. (In Press).

Duong, M. H. and Han, T. A. (2021b). Statistics of the number of equilibria in random social dilemma evolutionary games with mutation. *European Physical Journal B*. (In Press).

Endriss, U., Kraus, S., Lang, J., and Wooldridge, M. (2011). Incentive engineering for boolean games. *IJCAI '11*, pages 2602–2607.

Fehr, E. and Schmidt, K. M. (1999). A theory of fairness, competition, and cooperation. *The quarterly journal of economics*, 114(3):817–868.

Gut, A. (2009). *An Intermediate Course in Probability*. Springer Publishing Company, Incorporated, 2nd edition.

Güth, W. and Kocher, M. G. (2014). More than thirty years of ultimatum bargaining experiments: Motives, variations, and a survey of the recent literature. *Journal of Economic Behavior and Organization*, 108:396–409.

Güth, W., Schmittberger, R., and Schwarze, B. (1982). An experimental analysis of ultimatum bargaining. *Journal of economic behavior & organization*, 3(4):367–388.

Hamilton, W. (1964). The genetical evolution of social behaviour. i. *Journal of Theoretical Biology*, 7(1):1 – 16.

Han, T. A. (2013). *Intention Recognition, Commitments and Their Roles in the Evolution of Cooperation: From Artificial Intelligence Techniques to Evolutionary Game Theory Models*, volume 9. Springer SAPERE series.

Han, T. A., Lynch, S., Tran-Thanh, L., and Santos, F. C. (2018). Fostering cooperation in structured populations through local and global interference strategies. In *IJCAI-ECAI'2018*, pages 289–295.

Han, T. A., Pereira, L. M., and Santos, F. C. (2012). The emergence of commitments and cooperation. In *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems (AAMAS'2012)*, pages 559–566. ACM.

Han, T. A., Pereira, L. M., Santos, F. C., and Lenaerts, T. (2020). To Regulate or Not: A Social Dynamics Analysis of an Idealised AI Race. *Journal of Artificial Intelligence Research*, 69:881–921.

Han, T. A., Perret, C., and Powers, S. T. (2021). When to (or not to) trust intelligent machines: Insights from an evolutionary game theory analysis of trust in repeated games. *Cognitive Systems Research*, 68:111–124.

Han, T. A. and Tran-Thanh, L. (2018). Cost-effective external interference for promoting the evolution of cooperation. *Scientific reports*, 8(1):1–9.

Hutchinson, B. and Mitchell, M. (2019). 50 years of test (un)fairness: Lessons for machine learning. In *Proceedings of the Conference on Fairness, Accountability, and Transparency*, FAT* '19, page 49–58, New York, NY, USA. Association for Computing Machinery.

Lerat, J.-S., Han, T. A., and Lenaerts, T. (2013). Evolution of common-pool resources and social welfare in structured populations. In *IJCAI'2013*, pages 2848–2854.

Liu, L. and Chen, X. (2020). Evolutionary game dynamics in multiagent systems with prosocial and antisocial exclusion strategies. *Knowledge-Based Systems*, 188:104835.

Marton-Lafevre, J. et al. (2007). *Peace parks: conservation and conflict resolution*. Mit Press.

Maynard-Smith, J. (1982). *Evolution and the Theory of Games*. Cambridge University Press, Cambridge.

McAvoy, A. and Hauert, C. (2015). Asymmetric evolutionary games. *PLoS Comput Biol*, 11(8):e1004349.

Mehrabi, N., Morstatter, F., Saxena, N., Lerman, K., and Galstyan, A. (2021). A survey on bias and fairness in machine learning. *ACM Computing Surveys (CSUR)*, 54(6):1–35.

Nowak, M. A. (2006). *Evolutionary Dynamics: Exploring the Equations of Life*. Harvard University Press, Cambridge, MA.

Nowak, M. A. (2012). Evolving cooperation. *Journal of Theoretical Biology*, 299:1–8.

Nowak, M. A. and May, R. M. (1992). Evolutionary games and spatial chaos. *Nature*, 359(6398):826–829.

Nowak, M. A., Page, K. M., and Sigmund, K. (2000). Fairness versus reason in the ultimatum game. *Science*, 289(5485):1773–1775.

Ogbo, N. B., Elgarig, A., and Han, T. A. (2021). Evolution of coordination in pairwise and multi-player interactions via prior commitments. *Adaptive Behavior (In Press)*. Preprint arXiv:2009.11727.

Ohtsuki, H., Hauert, C., Lieberman, E., and Nowak, M. A. (2006). A simple rule for the evolution of cooperation on graphs and social networks. *Nature*, 441(7092):502–505.

Okada, I. (2020). A review of theoretical studies on indirect reciprocity. *Games*, 11(3):27.

Page, K. M., Nowak, M. A., and Sigmund, K. (2000). The spatial ultimatum game. *Proceedings of the Royal Society of London. Series B: Biological Sciences*, 267(1458):2177–2182.

Paiva, A., Santos, F. P., and Santos, F. C. (2018). Engineering pro-sociality with autonomous agents. In *Thirty-second AAAI conference on artificial intelligence*.

Perc, M., Jordan, J. J., Rand, D. G., Wang, Z., Boccaletti, S., and Szolnoki, A. (2017). Statistical physics of human cooperation. *Physics Reports*, 687:1–51.

Perc, M. and Szolnoki, A. (2010). Coevolutionary games—a mini review. *BioSystems*, 99(2):109–125.

Rand, D. G., Nowak, M. A., Fowler, J. H., and Christakis, N. A. (2014). Static network structure can stabilize human cooperation. *Proc Natl Acad Sci USA*, 111(48):17093–17098.

Rand, D. G., Tarnita, C. E., Ohtsuki, H., and Nowak, M. A. (2013). Evolution of fairness in the one-shot anonymous ultimatum game. *Proceedings of the National Academy of Sciences*, 110(7):2581–2586.

Rauwolf, P. and Bryson, J. J. (2018). Expectations of fairness and trust co-evolve in environments of partial information. *Dynamic Games and Applications*, 8(4):891–917.

Riehl, J. R. and Cao, M. (2016). Towards optimal control of evolutionary games on networks. *IEEE Transactions on Automatic Control*, 62(1):458–462.

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 the AAAI Conference on Artificial Intelligence*, volume 33, pages 6146–6153.

Sasaki, T., Brännström, Å., Dieckmann, U., and Sigmund, K. (2012). The take-it-or-leave-it option allows small penalties to overcome social dilemmas. *Proceedings of the National Academy of Sciences*, 109(4):1165–1169.

Selten, R. (1978). A note on evolutionarily stable strategies in asymmetric animal conflicts.

Sigmund, K. (2010). *The Calculus of Selfishness*. Princeton University Press.

Sinatra, R., Iranzo, J., Gomez-Gardenes, J., Floria, L. M., Latora, V., and Moreno, Y. (2009). The ultimatum game in complex networks. *Journal of Statistical Mechanics: Theory and Experiment*, 2009(09):P09012.

Smidt, H. M. (2020). United nations peacekeeping locally: enabling conflict resolution, reducing communal violence. *Journal of Conflict Resolution*, 64(2-3):344–372.

Szabó, G. and Fath, G. (2007). Evolutionary games on graphs. *Physics reports*, 446(4-6):97–216.

Szabó, G. and Fáth, G. (2007). Evolutionary games on graphs. *Phys Rep*, 97-216(4-6).

Teixeira, A. S., Santos, F. C., Francisco, A. P., and Santos, F. P. (2021). Eliciting fairness in n-player network games through degree-based role assignment. *Complexity*, 2021.

Traulsen, A., Hauert, C., De Silva, H., Nowak, M. A., and Sigmund, K. (2009). Exploration dynamics in evolutionary games. *Proceedings of the National Academy of Sciences*, 106(3):709–712.

Traulsen, A., Nowak, M. A., and Pacheco, J. M. (2006). Stochastic dynamics of invasion and fixation. *Phys Rev E*, 74(1):011909.

Trivers, R. L. (1971). The evolution of reciprocal altruism. *Quaterly Review of Biology*, 46:35–57.

Tuyls, K. and Parsons, S. (2007). What evolutionary game theory tells us about multiagent learning. *Artificial Intelligence*, 171(7):406–416.

Tuyls, K., Perolat, J., Lanctot, M., Ostrovski, G., Savani, R., Leibo, J. Z., Ord, T., Graepel, T., and Legg, S. (2018). Symmetric decomposition of asymmetric games. *Scientific Reports*, 8(1):1–20.

Wang, S., Chen, X., and Szolnoki, A. (2019). Exploring optimal institutional incentives for public cooperation. *Communications in Nonlinear Science and Numerical Simulation*, 79:104914.

West, S., Griffin, A., and Gardner, A. (2007). Evolutionary explanations for cooperation. *Current Biology*, 17:R661–R672.

Wilder, B., Immorlica, N., Rice, E., and Tambe, M. (2018). Maximizing influence in an unknown social network. In *AAAI conference on Artificial Intelligence (AAAI-18)*.

Wooldridge, M. (2012). Bad equilibria (and what to do about them). *ECAI '12*, pages 6–11.

Zisis, I., Guida, S. D., Han, T. A., Kirchsteiger, G., and Lenaerts, T. (2015). Generosity motivated by acceptance - evolutionary analysis of an anticipation games. *Scientific reports*, 5(18076).

5 Supplementary Information

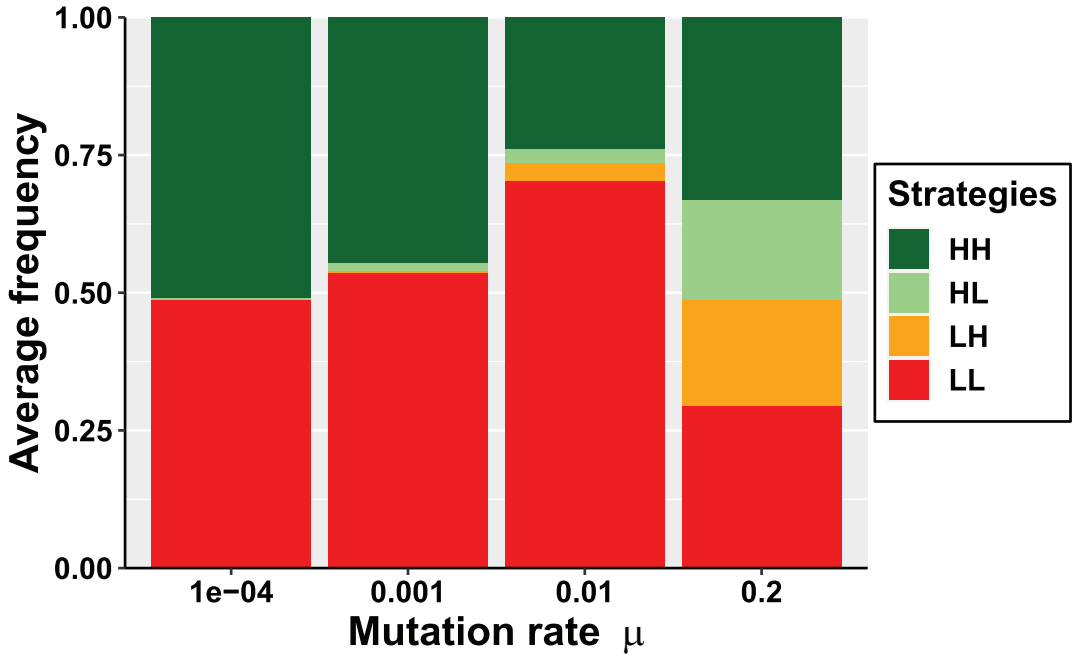

Figure S1. Average frequencies of the four strategies HH, HL, LH and LL as a function of mutation rate μ in absence of interference.

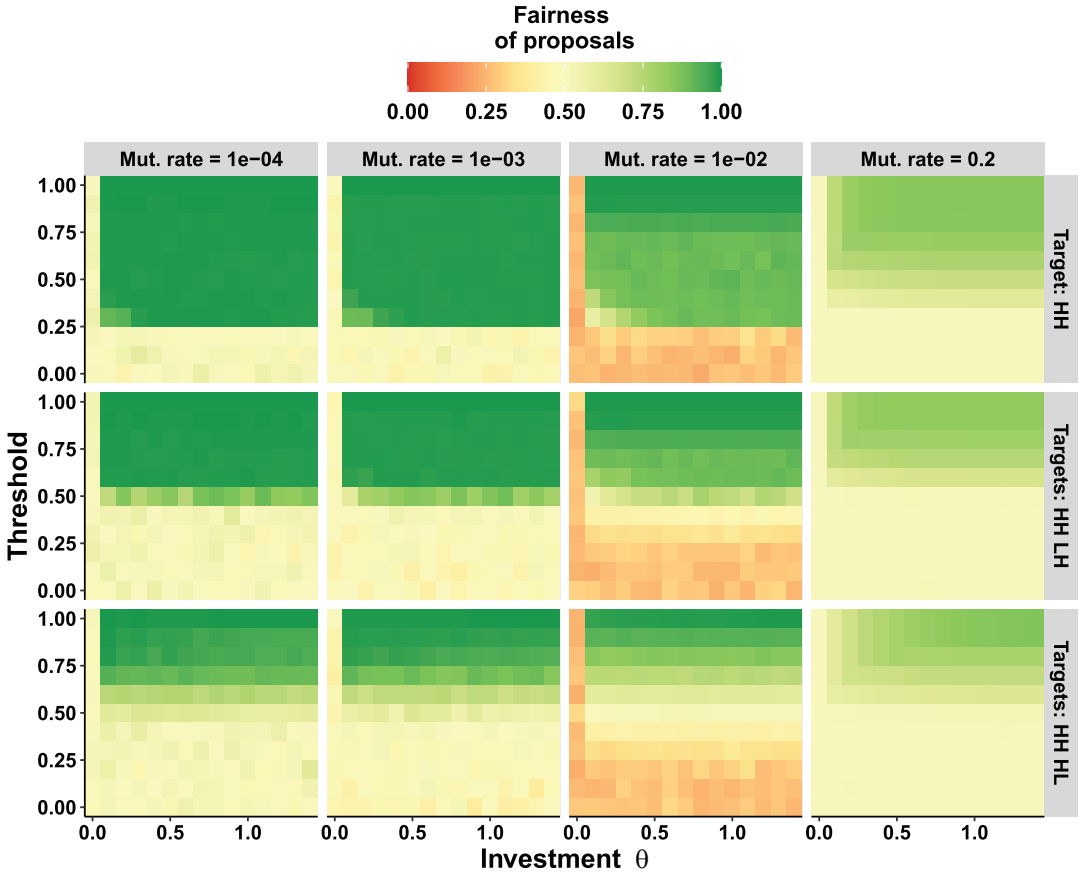

Figure S2. Average fairness as a function of the individual endowment θ, the threshold p_f and the mutation rate μ (population-based, stochastic update). Each row represents a different targeting scheme.

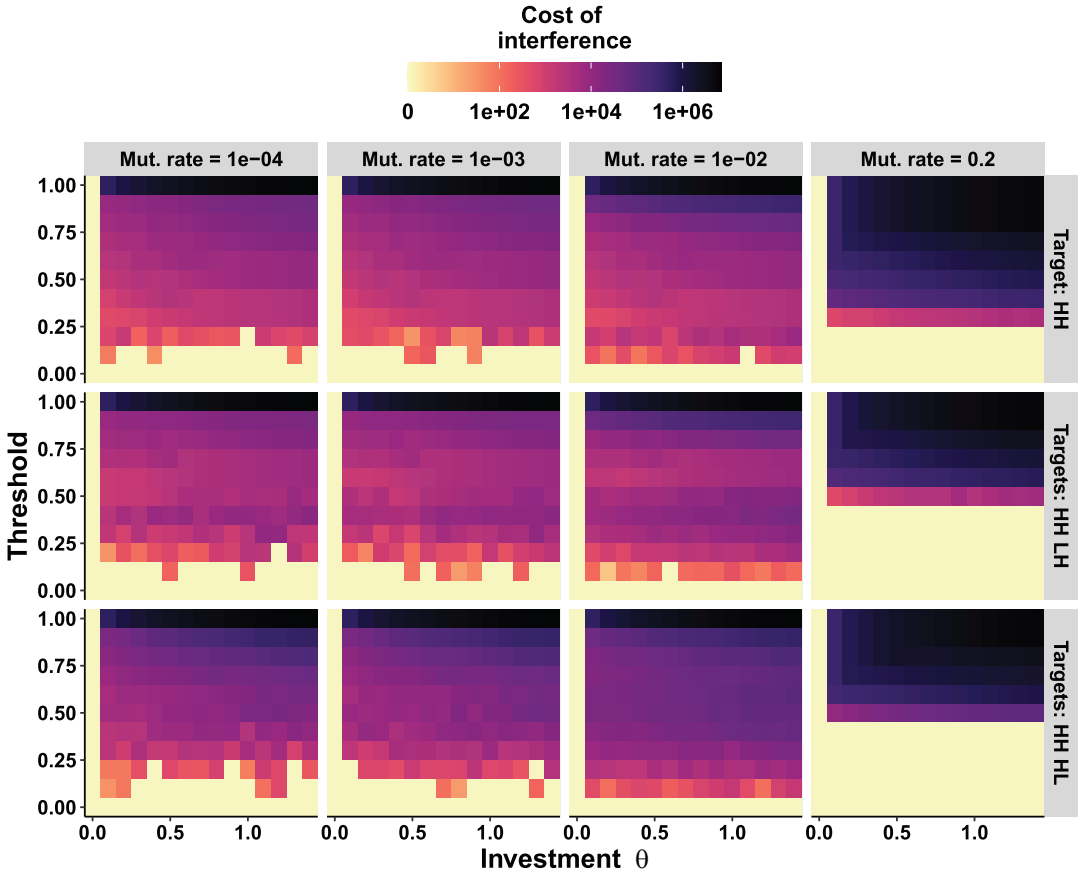

Figure S3. Average cost of interference as a function of the individual endowment θ, the threshold p_f and the mutation rate μ (population-based, stochastic update). Each row represents a different targeting scheme. The cost of interference is shown on a logarithmic scale.

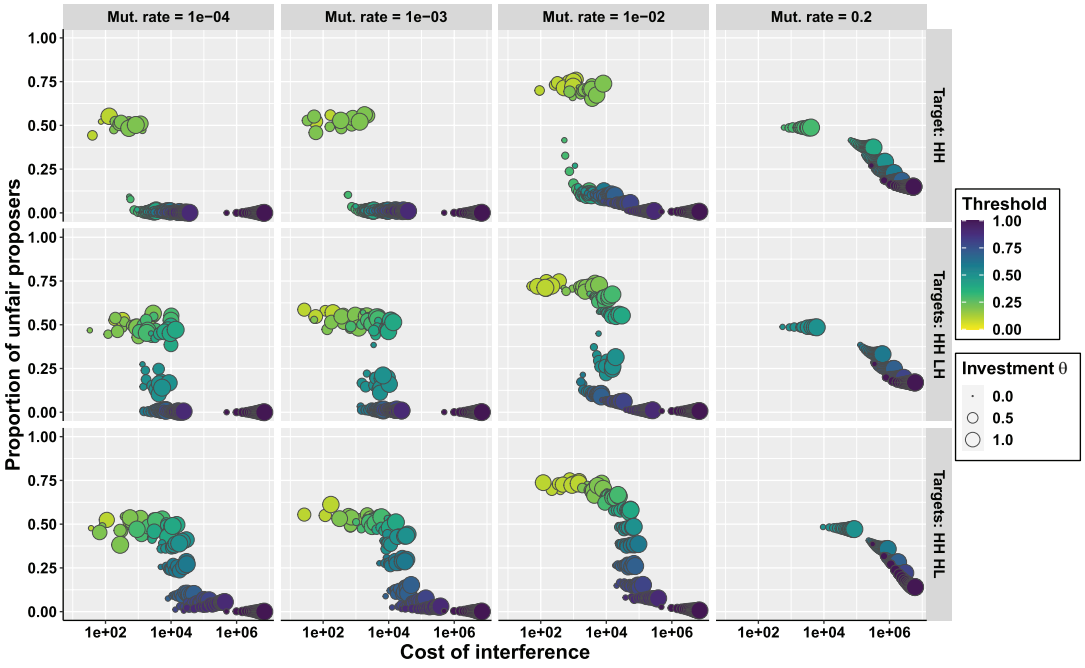

Figure S4. Proportion of unfair proposers as a function of average cost of interference for different targeting scheme and mutation rate μ (population-based, stochastic update). The size and colour of the circles correspond to investment amount and threshold of investment, respectively. We note that the most desirable outcomes are closest to the origin.

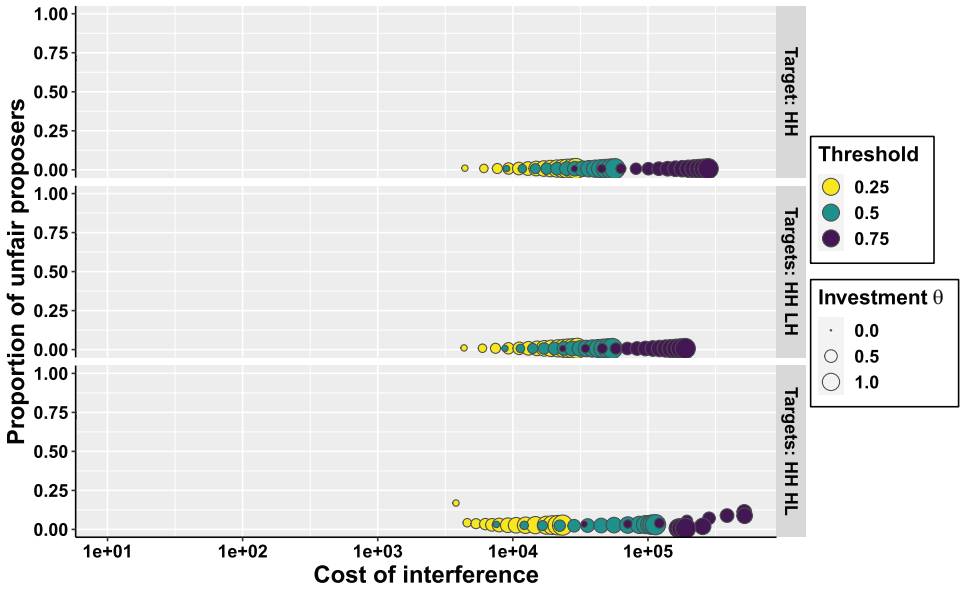

Figure S5. Proportion of unfair proposers as a function of average cost of interference for different targeting scheme (neighbourhood-based, $\mu = 0.01$, stochastic update). The size and colour of the circles correspond to investment amount and threshold of investment, respectively. We note that the most desirable outcomes are closest to the origin.

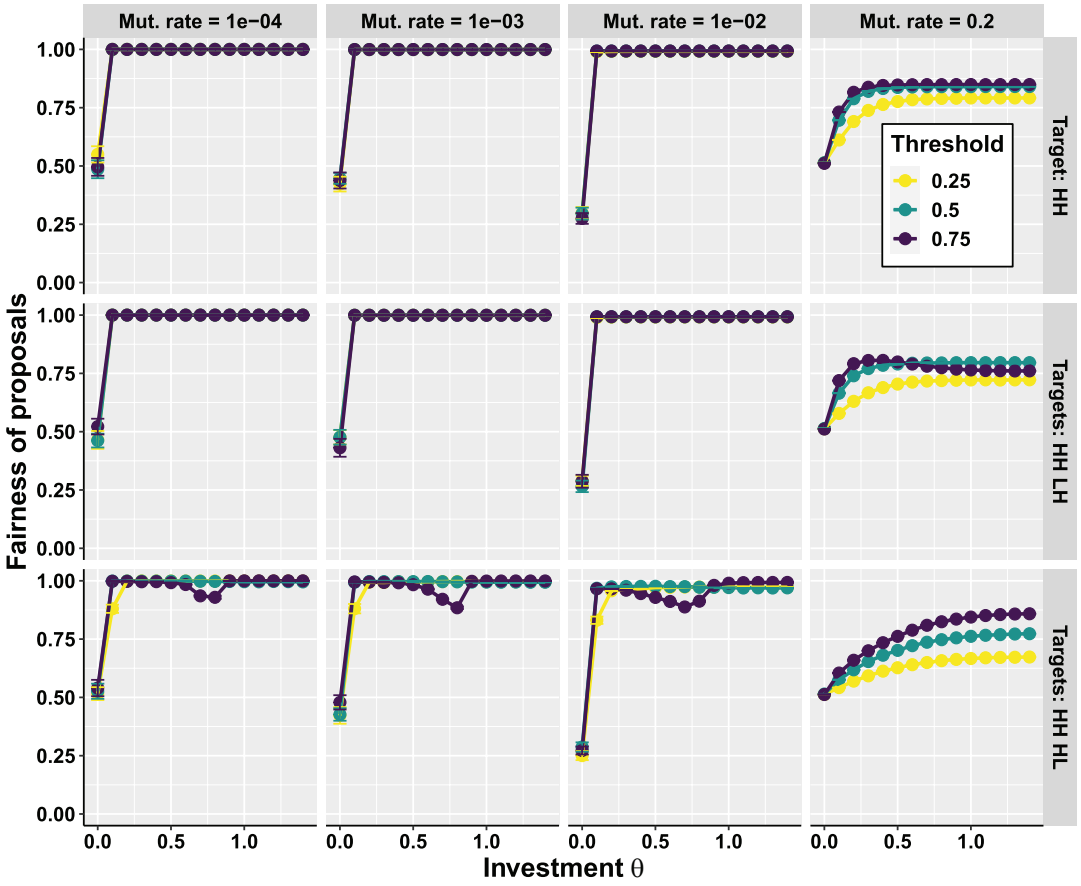

Figure S6. Average fairness measured by the sum of frequencies of HH and HL as a function of the individual endowment θ, the threshold p_f and the mutation rate μ (neighbourhood-based, stochastic update). Each row represents a different targeting scheme.

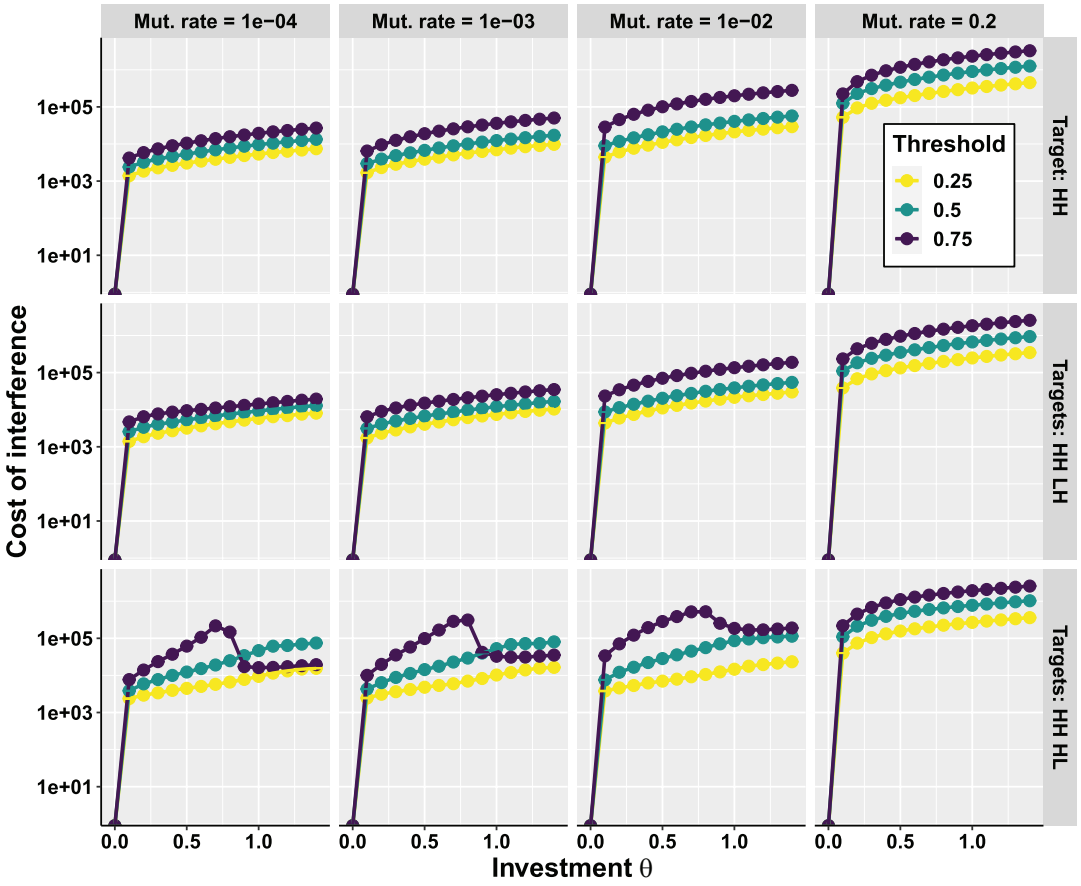

Figure S7. Average cost of interference as a function of the individual endowment θ, the threshold p_f and the mutation rate μ (neighbourhood-based, stochastic update). Each row represents a different targeting scheme. The cost of interference is on a logarithmic scale for clarity.

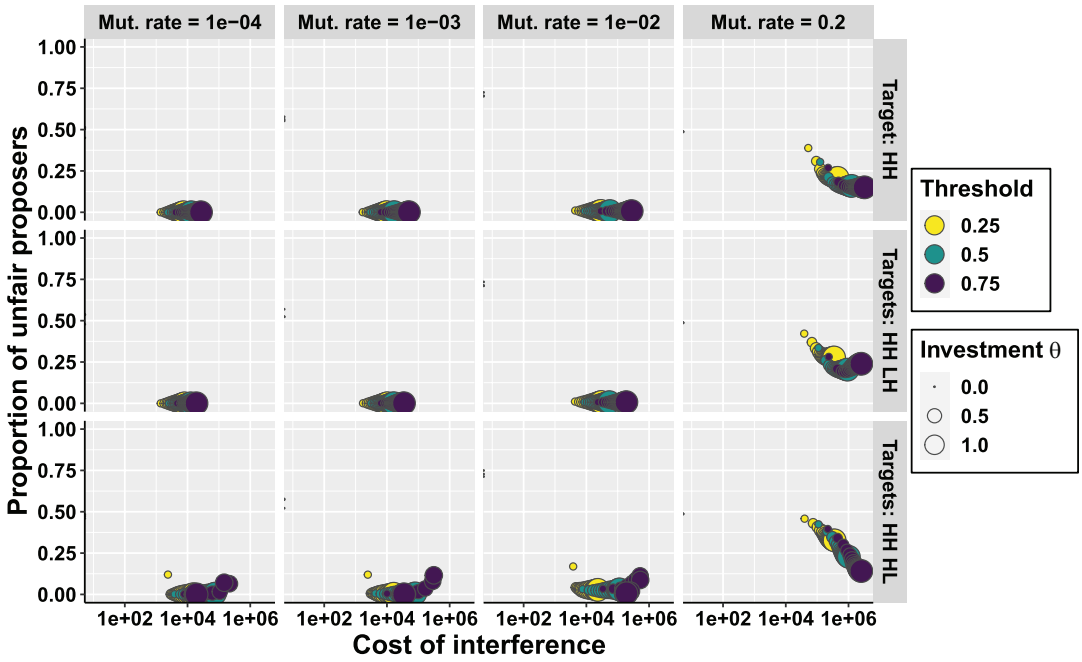

Figure S8. Proportion of unfair proposers as a function of average cost of interference for different targeting scheme and mutation rate μ (neighbourhood-based, stochastic update). The size and colour of the circles correspond to investment amount and threshold of investment, respectively. We note that the most desirable outcomes are closest to the origin.

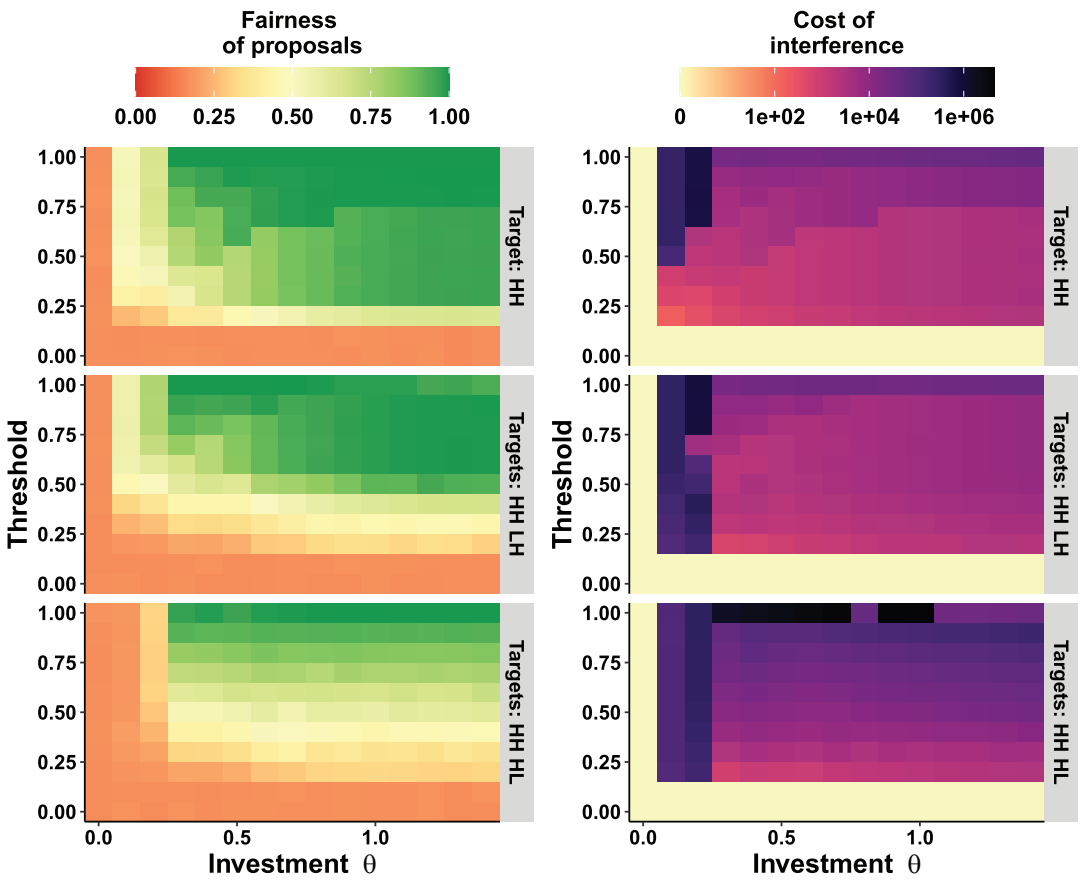

Figure S9. Average fairness (left) and average cost of interference (right) as a function of the individual endowment θ and the threshold p_f (population-based, deterministic update). Each row represents a different targeting scheme. The cost of interference is shown on a logarithmic scale.