

Effects of parity, sympathy and reciprocity in increasing social welfare

SANDIP SEN[✉], CHAD CRAWFORD, ADAM DEES, RACHNA NANDA KUMAR, and JAMES HALE

The University of Tulsa

e-mails: sandip@utulsa.edu, chad-crawford@utulsa.edu, apd615@utulsa.edu, rachmanandakumar@utulsa.edu, jah6484@utulsa.edu

Abstract

We are interested in understanding how socially desirable traits like sympathy, reciprocity, and fairness can survive in environments that include aggressive and exploitative agents. Social scientists have long theorized about ingrained motivational factors as explanations for departures from self-seeking behaviors by human subjects. Some of these factors, namely reciprocity, have also been studied extensively in the context of agent systems as tools for promoting cooperation and improving social welfare in stable societies. In this paper, we evaluate how other factors like sympathy and parity can be used by agents to seek out cooperation possibilities while avoiding exploitation traps in more dynamic societies. We evaluate the relative effectiveness of agents influenced by different social considerations when they can change who they interact with in their environment using both an experimental framework and a predictive analysis. Such rewiring of social networks not only allows possibly vulnerable agents to avoid exploitation but also allows them to form gainful coalitions to leverage mutually beneficial cooperation, thereby significantly increasing social welfare.

1 Introduction

The goal of a rational agent is to maximize the utility received from interactions with its environment. In single-agent systems, this assumption leads to choosing actions that maximize expected utility. Even in single-agent scenarios, though, one has to be careful in differentiating between greedy choices, that maximizes short-term gains or improvements, versus more strategic action choice mechanisms that seeks to maximize longer-term, albeit discounted, utilities. The latter approach is preferred as it improves overall viability and success of the agent, even though it means less immediate benefits, and even short-term losses in some cases. In a multiagent context, it is not useful to seek unilateral benefits in the presence of other agents. Early results in game theory showed that to guarantee safety values in multistage games, one has to adopt minimax strategies that take into consideration the desires of other agents to maximize their payoff. Concomitantly, a large body of literature in simultaneous move, single-stage games has studied human behaviors motivated by altruism, reciprocity, etc. While social scientists have developed theories about why such behavior is prevalent in human societies, agent researchers have tried to identify effects of similar considerations in enabling and sustaining cooperative relationships in agent societies. This paper studies three motivational factors that suggest a clear departure from self-utility maximization goals that have been identified by social scientists to be influential in human decision making. These commonly observed factors are:

Sympathy: In addition to consideration of self-utility, individuals take into consideration the utility received by others and can seek to benefit others even at the expense of local cost (Dawes & Thaler, 1988).

Parity: Individuals are additionally prone, in many situations, to hold equitable outcomes in high esteem, and prefer such outcomes over those that would lead to higher local utility (Bolton, 1991; Bolton & Ockenfels, 2000).

Reciprocity: While interacting with a particular partner, individuals are at times motivated to return favors and slights, that is, helping gestures are reciprocated and hurtful actions are penalized (Goranson & Berkowitz, 1966; Martinez-Coll & Hirshleifer, 1991; Rabin, 1993; Fehr *et al.*, 1997).

The above behavioral traits make sense in gregarious human societies: we live in groups and communities. Relationships are at least semi-stable and involve repeated interactions. Reputation and trust are key social capitals that can protect us or inform our decisions when we are at a bind or meeting new acquaintances. Various evolutionary forces, including kin selection, as well as egotistical reasons ('I would like to be seen as the good guy'), can motivate us to deviate from purely self-interested behavior, even without guaranteed long-term returns for not maximizing self-utility in each local interaction. When applied to agent societies, such traits can be incorporated in agent designs to reflect the preferences and biases of their human counterparts. But do these traits add to the competitiveness of agents?

We first experimentally investigate this perplexing question in the context of agents repeatedly interacting with neighboring agents located on a small-world (Watts & Strogatz, 1998) social network. Agent interactions are in the form of stylized games which allow the agents to adopt from simple choices of cooperation and defection, with deterministic payoffs for each possible outcome known *a priori* to all parties. Agents can recall interaction outcomes with past partners and current neighbors. In addition, agents can choose to sever ties with neighbors from whom they receive unsatisfactory interaction utilities and connect with others with whom interactions are expected to be more rewarding. We populate such societies with both purely self-utility-maximizing agents as well as agents whose decision making is influenced by any one or all of the three motivational factors of sympathy, parity, and reciprocity. We observe the cumulative payoff (welfare) received by various agent types over a number of interactions and analyze the evolving topology of the network of connections between the agents. We experiment with various heterogeneous agent groups to identify the relative superiority of agent types against each other as well as how they perform when all agent types co-habituate. A number of non-intuitive, yet telling, details emerge from these set of experiments: (a) the head-to-head dominance pattern of the agents reveal a cyclic pattern, (b) pure selfishness is self-defeating, (c) whereas sympathy is useful in increasing social and individual welfare in many situations, parity is rarely useful, (d) how influential a motivational factor is on one's decision making determines dominant behaviors together with the initial agent type distribution, (e) agents who are influenced by multiple motivational factors may outperform others who are motivated by a single factor, and (f) the ability to rewire one's social connections is key to the viability and vibrancy of a dynamic society.

We also present an analysis of expected utilities in heterogeneous populations of different agent types to predict the expected total payoffs of strategy types. We use the analysis to identify equilibrium configurations and determine stable partnerships. Such a model enables us to see in what initial conditions certain behaviors thrive in, thus providing us the knowledge and tools to construct systems which favor the desirable qualities such as high social welfare or fairness.

The rest of the paper is organized as follows: Section 2 discusses some pertinent literature, Section 3 presents the environmental and agent characteristics used in this study, Section 4 presents results from various experimental runs with different population configurations, Section 5 presents a predictive analysis of heterogeneous agent populations to identify the effectiveness of different strategies from a given starting population type distributions over a small-world agent network, and Section 6 summarizes the takeaways from this study and identifies potentially fruitful research directions.

2 Related work

Given the rapid increase in user interest and participation in online social networks, researchers are focusing on understanding how interactions between individuals lead to emergent social structures and phenomena (David & Jon, 2010; Baetz, 2015; Barabasi, 2016), such as how individuals influence and are

influenced by other users they are connected to Cha *et al.* (2010). Other researchers have used agent-based models and simulations to explore how behavioral traits and strategic interaction decisions can influence social network dynamics such as change in topologies (Galán *et al.*, 2011), information flow (Tsang & Larson, 2014), or to characterize the emergence of conventions or norms (Epstein, 2001; Delgado, 2002) or cooperative behavior (Mahmoud *et al.*, 2016). Some of these projects are formal studies to prove convergence derive rational agent behaviors (Brooks *et al.*, 2011). Others utilize extensive experimental evaluations to characterize the nature of emerging behaviors and topologies in networks of self-interested agents (Delgado, 2002; Peleteiro *et al.*, 2014; Mahmoud *et al.*, 2016).

Some of these research investigate how the network topology changes based on strategic or exploratory rewiring of connections by agents seeking relationships that produce higher payoffs (Peleteiro *et al.*, 2014). Interaction between network neighbors are often represented as a stage game (Epstein, 2001; Delgado, 2002), for example, in studies on the emergence of conventions, the chosen stage games are symmetric coordination games with multiple Nash equilibria.

The behavioral traits of sympathy, parity and reciprocity have been inspired by a study involving human subjects (Güth & Yaari, 2004) on the Prisoners Dilemma game (Roth, 1995). The basic premise of the paper is that people do not necessarily try to maximize their own payoff from interactions with others. Other motivational forces, that take into consideration the payoffs received by interaction partners, have been repeatedly observed to have a key influence on human decision making. As a result, researchers have found notable departures of human actors from self-seeking behaviors (Dawes & Thaler, 1988; Sally, 1995). While some researchers have observed considerations of *parity* or *inequality aversion* when a user is in a disadvantageous position, others have found its use in situations without power imbalance (Bolton, 1991; Bolton & Ockenfels, 2000). Sympathy, which suggests the ability to put oneself in the position of one's partner, has been traditionally used to explain cooperative behavior that cannot be rationalized by self-interest (Dawes & Thaler, 1988). Social scientists have devoted considerable attention to a somewhat extreme and paradoxical form of sympathy, *altruism*, where benefits to others is given exclusive consideration (Krebs, 1970; Taylor, 1992; Badhwar, 1993; Schmitz, 1993; Day & Taylor, 1998). Reciprocity (Goranson & Berkowitz, 1966), that is, the desire to do to others as they do to you or applying *measure for measure*, has also been recognized as a driving force behind human decision making (Martinez-Coll & Hirshleifer, 1991; Rabin, 1993; Fehr *et al.*, 1997). For example, experiments with human subjects on sequential social dilemma games show that subjects will incur personal costs to penalize norm violators or reward norm upholders (Fehr & Fischbacher, 2003). Additionally, some of the experiments are carried out in public goods game scenarios. In contrast, we study simultaneous move game in societies without the background of existing norms and interactions produce private payoffs only. Others have used centralized monitoring and penalty mechanisms to ensure fair access to resources in a society (Danassis & Faltings, 2019). We do not assume the presence of any mediating agents or capability to impose direct monetary punishment for individuals for aberrant behavior.

Researchers have looked at local Nash Equilibrium in structured populations, for example, social networks, and the resultant effects on the distribution of defectors and cooperators in social dilemma games such as the Prisoner's dilemma as well as the evolution of stable strategies (Zhang *et al.*, 2014). In contrast to this work, our agents can change their connections but do not change their strategies.

Güth & Yaari (2004) attempt to explain human considerations of sympathy, parity, and reciprocity when departing from self-seeking behavior. In contrast, they try to characterize scenarios where consideration of partner payoffs can maximize utility for self-seeking autonomous agents, trying to maximize cumulative interaction payoffs from their partners. Social networks where an agent's utility depended on their neighbors' payoffs have been studied in the context of social choice (Salehi-Abari & Boutilier, 2014). The focus of this work is on better social choice mechanisms rather than the viability of individual agent types as is the case in our work. Various computational approaches have studied a certain form of reciprocity, mimicking an opponent's past action, in promoting cooperation in repeated game scenarios and in multiagent societies (de Vos & Zeggelink, 1994; Ferriere & Michod, 1996; Sen, 1996). In this paper, a more strategic form of reciprocity is used that takes into consideration partner payoffs to determine strategy. More specifically, a unified computational framework for decision making, that

Table 1 Raw Payoffs in a Prisoner's Dilemma game
(default values used (Güth & Yaari, 2004): $\alpha = 5$, $\epsilon = 0.1$)

	C	D
C	(α, α)	$(-\epsilon, \alpha + \epsilon)$
D	$(\alpha + \epsilon, -\epsilon)$	$(0,0)$

Table 2 Raw Payoffs in a battle of sexes game

	C	D
C	$(2,1)$	$(0,0)$
D	$(0,0)$	$(1,2)$

incorporates sympathy, parity, and reciprocity considerations, together with payoff to self, is used to select actions. To the best of our knowledge, such unified treatment has not been attempted before.

3 Model

This section describes the game scenario used in this paper to represent the interaction between two players, the player types, and strategic decision making framework, including actions for severing and forming social network connections. The agents (players) are situated on a social network, where each agent corresponds to a node and edges connect agents to other agents they can interact (play) with.

To initialize, a Watts-Strogatz model (Watts & Strogatz, 1998) is used to generate a *small world network* with agents allocated to randomly selected nodes. A simulation run consists of repeated rounds numbering R . Every round, each agent is called upon to initiate an interaction (a game) with another agent, which it selects uniformly from all of its neighbors. When an agent initiates an interaction, it will be the first to make its strategy choice. The players move in turns: in a sequential two-player, two-action game, the first player's strategy choice is visible to the second player when the latter chooses its action. An agent is not exempt from being part of another interaction in the same round should a different agent initiate an interaction upon it (although it will be the second player in these instances).

In addition, after R_r rounds, each interaction phase will be followed by a rewiring phase. The rewiring phase grants each agent two opportunities: to form and sever a connection with other agents in the network. Agents connect to previously unknown agents via the recommendation of their most favored partner, and disconnect from agents based on a global connection cost constant Γ .

3.1 Game

The games used in this paper are based on sequential variants of the Prisoner's Dilemma and Battle of Sexes games. The Prisoner's Dilemma payoff matrix for two interacting agents is presented in Table 1, where $\alpha > 0$, $\epsilon > 0$ (we use $\alpha \gg \epsilon$ following Güth & Yaari (2004)). The first value of each pair, x , is the row (first) player's pay off, and the second value, y , is the column (second) player's pay off. Similarly, the payoff matrix for the Battle of Sexes is presented in Table 2.

The utility to an agent for a particular outcome depends on the payoff of both players as the players are not only self-seeking but are also influenced by sympathy, parity, and reciprocity considerations. The utility of a player who receives a payoff x when the other player receives a payoff of y from an interaction is calculated as

$$u(x, y) = w_{me}x + w_s y - w_p |x - y| + w_r (C\rho_c - (1 - C)\rho_d)y$$

Table 3 Weight vectors of each of the four agent types

	w_s	w_p	w_r	w_{me}	ρ_c	ρ_d
Selfish	0	0	0	1	–	–
Sympathy	0.5	0	0	0.5	–	–
Parity	0	0.5	0	0.5	–	–
Strict reciprocative	0	0	0.5	0.5	1	1
Considerate reciprocative	0	0	0.5	0.5	1	0.5
Mixed	0.25	0.25	0.25	0.25	1	1

Table 4 Utility to different player types for all possible Prisoner’s Dilemma game outcomes (with the first and second letters representing the strategy choices of that row’s type and its partner respectively)

	CC	CD	DC	DD
Selfish	5	–0.1	5.1	0
Sympathy	5	2.5	2.5	0
Parity	2.5	–2.65	–0.05	0
Strict reciprocative	5	–2.6	2.5	0
Considerate reciprocative	5	–1.325	2.5	0
Mixed	3.75	–1.325	–0.075	0

where C is 1 or 0 depending on if the other player cooperates or defects respectively, ρ_c and ρ_d are the fraction of the other player’s payoff used by a reciprocative player when the opponent cooperates and defects respectively.

We define an agent type, τ , as a vector of six parameters:

$$\tau = \langle w_s^\tau, w_p^\tau, w_r^\tau, w_{me}^\tau, \rho_c, \rho_d \rangle$$

where the first four parameters represent the relative influences of sympathy, parity, reciprocity and selfishness, respectively, on the agent, and the last two deal with weights toward reciprocating cooperation and defection respectively, if an agent type makes use of reciprocity. The weights normalize the influences on utility, that is, $w_s + w_p + w_r + w_{me} = 1$. We experimented with six types of agents; the corresponding weight vectors are presented in Table 3. The Utility matrix for each of these player types for each of the outcomes is presented in Table 4 for the Prisoner’s Dilemma game. An outcome is characterized by the strategy combination played, where the first letter is the strategy used by the current agent and the second letter is the strategy chosen by its opponent.

3.2 Types of players

Selfish: A player whose utility is equal to its own payoff: $w_{me} = 1$ and all other factors are assigned 0 weight.

Sympathy: A player who has utility equally split between their own payoff and that of their opponent¹. We use the agent type $\langle 0.5, 0, 0, 0.5, \rho_c, \rho_d \rangle$.

Parity: A player that not only considers its own payoff but also tries to minimize the inequality in the payoffs between the two players: given by weight vector $\langle 0, 0.5, 0, 0.5, \rho_c, \rho_d \rangle$.

¹ In literature, often an altruistic agent is discussed (Krebs, 1970; Taylor, 1992; Badhwar, 1993; Schmitz, 1993; Day & Taylor, 1998), which can be viewed as an extreme case of sympathetic agent which ignores its own payoff and puts the entire weight on the opponent’s payoff when calculating its utility of an outcome, that is, for an altruist, $w_{me} = 0$, $w_s = 1$, and its type is described by the weight vector $\langle 1, 0, 0, 0, -, - \rangle$. The sympathetic agents we use balances selfish considerations with consideration about the well-being of the opponent.

Reciprocity: A player that includes reciprocal consideration of the opponent's choice to determine its utility: given by weight vector $\langle 0, 0, 0.5, 0.5, \rho_c, \rho_d \rangle$. If the opponent cooperates (resp. defects), the reciprocal player adds (resp. subtracts) $w_r \rho_c$ (resp. $w_r \rho_d$), of the other player's payoff to its utility.

We use two types of reciprocal players:

Strict Reciprocal: Agents add or subtract the same percentage of the other player's pay off to compute its utility (we use $\rho_c = \rho_d = 1$), and

Considerate Reciprocal: Agents are less harsh to the opponent: given a uniform prior over opponent strategies, the player chooses cooperation over defection. Now for a reciprocity agent

$$EU[C] = 0.25(\alpha(1 + \rho_c - \rho_d) - \epsilon(1 + \rho_d)),$$

$$EU[D] = 0.25(\alpha + \epsilon(1 - \rho_c))$$

For a considerate reciprocal player, we desire $EU[C] \geq EU[D]$, that is, $\rho_c - \rho_d > \frac{2\epsilon}{\alpha + \epsilon}$. For $\alpha = 5$, $\epsilon = 0.1$, $\rho_c - \rho_d \geq 0.039$. We select $p_c = 1$ and $p_d = 0.5$ to force first move cooperation for the considerate reciprocal type.

Mixed: Finally, we envision a player who assigns equal weight to all influence factors: their own payoff, the other player's payoff, the difference in the two players' payoffs, and the other player's strategy. Hence, the mixed player type is described by the vector $\langle 0.25, 0.25, 0.25, 0.25, 1, 1 \rangle$.

3.3 Strategy decision

Every time a player interacts with another player, it records the other player's strategy and uses this historical record the next time it interacts with the same player to estimate the probability the other player cooperates or defects. Suppose, an opponent is observed to have cooperated t_c times and defected t_d times with this player. Then the next interaction with this player, it assumes that the player will cooperate with a probability, $P_c = t_c / (t_c + t_d)$ and defects with a probability, $P_d = t_d / (t_c + t_d)$. In the special case of no knowledge (i.e., the very first move), $P_c = P_d = 0.5$ is assumed.

The strategy chosen by a player, when making the first move, depends on its expected utility value of cooperating, $EU[C]$, and defecting, $EU[D]$, given the expected move of its opponent, and is calculated as follows:

$$EU[C] = P_c u(CC) + P_d u(CD)$$

$$EU[D] = P_c u(DC) + P_d u(DD)$$

If the player is the second player, then its expected utility for its actions is calculated from its utility value for corresponding outcomes because the opponent's strategy is already known. So the utility for the second player in the first interaction is calculated as $EU[C] = u(CC)$ or $EU[D] = u(DC)$ when the other player cooperates, and $EU[C] = u(CD)$ or $EU[D] = u(DD)$, when the other player defects. A player decides to cooperate when $EU[C] \geq EU[D]$ and defects otherwise. The raw expected utility that a player has for an interaction is $EU = \max\{EU[C], EU[D]\}$.

3.4 Network and connections

Except for the experiments on pairwise agent interactions, all other networks are initialized using the Watts-Strogatz network small-world generation algorithm (Watts & Strogatz, 1998). The Watts-Strogatz algorithm initially places agents on a ring-structured network where each node has degree n , and then randomly reconnects each edge on the network with probability p (we use $n = 4$ and $p = 0.01$). The parameters for the algorithm were chosen so that the network was well-connected for agents to have flexibility in choosing their neighbors, while still representing a realistic structure. A small-world network was ideal for these experiments since all agents initially have approximately equal influence. Only once

Table 5 Notation

Symbol	Meaning
\mathcal{A}	Set of agents
$type_i$	Set of agents of type i
(x,y)	Payoffs for agents 1 and 2 respectively
$u(x,y)$	Utility for agent 1 if its payoff is x and that of its opponent is y
τ	Agent type
w_s	Sympathetic weight
w_{me}	Selfish weight
w_p	Parity weight
w_r	Reciprocity weight
ρ_c	Reciprocate cooperation weight
ρ_d	Reciprocate defection weight
P_c	Prior defection proportion for other agent
P_d	Prior cooperation proportion for other agent
$EU[C]$	Expected utility for an agent assuming it cooperates
$EU[D]$	Expected utility for an agent assuming it defects
n	Mean degree
R	Total rounds
$R_{\bar{r}}$	No rewiring rounds
R_r	Rewiring rounds

the rewiring phase of the simulation begins (unrelated to Watts-Strogatz reconnection/rewiring), will the influences of the agents in the network begin to differ.

After $R_{\bar{r}} = 100$ iterations of the interaction stage of the algorithm, each agent is given the choice to connect to a new agent and disconnect from an undesirable, existing neighbor. The motivation behind this process is that by 100 iterations, the agents would have had time to interact with all the types of agents and realize which types of agents will be helpful and which types will hurt them in an interaction, and can thus seek new connections that could replace these ‘uncooperative’ individuals.

During the connection phase, an agent a requests a recommendation from its ‘preferred partner’ that is, the neighbor it has the highest expected utility for, $EU = \max\{EU[C], EU[D]\}$, as the first player in the game. This ‘highest expected utility’ is equivalent to a trust value for the agent (Gambetta, 1988; Sen, 2013): the higher the expected utility, the more the agent trusts their partner to play a strategy that is mutually rewarding. The preferred partner agent, b , recommends another agent c using a weighted sample based on expected utility (with negative EU partners discarded) from b that is not connected to a . b reports the expected utility it has for c , to a . Agent a will trust this estimate of c unless a has prior experience with c , in which case it will use the information from its prior experience to determine its own expected utility.

Following this recommendation stage, a elects to form a new connection with this neighbor if and only if the expected utility is higher than the connection cost parameter Γ . The connection cost is the cost of forming or maintaining social relationships on the network. It can be interpreted as a ‘cognitive load’ on the agent, such that interactions with others only pay dividends when the reward outweighs the cognitive cost.

Disconnecting works as follows: after the recommendation and connection stages, each agent is given the choice to disconnect from their ‘least preferred’ partner, d , which is their partner with the lowest expected utility, EU. If $EU \leq \Gamma$, the agent will disconnect from this partner. Otherwise, the agent will keep all existing connections.

We summarize the notations introduced herein for easy reference in Table 5.




Figure 1 Prisoner's Dilemma: Equilibrium strategies with (a) Strict Reciprocity and (b) Considerate Reciprocity agents. Edges originating from $type_i$ to $type_j$ represent the equilibrium utilities when $type_i$ plays the first move

4 Experimental results

Unless otherwise stated, all games have been simulated with 100 agents for $R = R_r + R_{\bar{r}} = 250$ iterations, allowing re-wiring only after $R_{\bar{r}} = 100$ iterations.

The results of the Prisoner's Dilemma game are presented in greater depth as the Battle of Sexes game was used primarily as a means to compare and contrast the results obtained with the Prisoner's Dilemma game to those in a different domain.

4.1 Pairwise interactions in one-on-one games

It is important to understand the convergence in players' strategies in repeated binary interactions, because those equilibrium strategy profiles will also be observed across the population when agents repeatedly interact with their neighbors on a social network. Convergence here refers to an equilibrium, understood as a state where every interaction afterward will produce the same outcome. The corresponding convergence strategy utilities for Strict reciprocity and Considerate reciprocity are shown in Figure 1. The convergence strategy profile for Considerate Reciprocity agents differed from the other case only in the following ways: Reciprocity to Parity is CC instead of DD and Reciprocity to Reciprocity is CC instead of DD. These two changes occur because considerate reciprocative opens with C when it goes first, and strict reciprocative opens with D: behaviors forced due to our intentional selection of ρ_d and ρ_c values.

4.2 Pairwise interactions in larger groups

To understand the nature of pairwise interactions on social networks, a group of 50 agents of one type are played against another group of 50 agents of another type. To observe which group (multiple players of one type) does better than the other, two groups of every type pair were played against each other and the average payoffs were recorded. The difference in this simulation from the two player simulation is that here not only can a player interact with a player of the other type but it can also interact with other players of its own type.

For the Prisoner's Dilemma game, these simulations showed that there can be cyclic patterns in the dominance order of agent types in terms of who received the higher payoff in these head-to-head matches. In all cases, the Mixed type dominates other agents influenced primarily by one social consideration as well as the selfish agents. Strict reciprocative players tied with parity and selfish players; all three types having an equal payoff when played against each other. However, parity and strict reciprocative agents performed worse than sympathy in pairwise interactions, while selfish agents

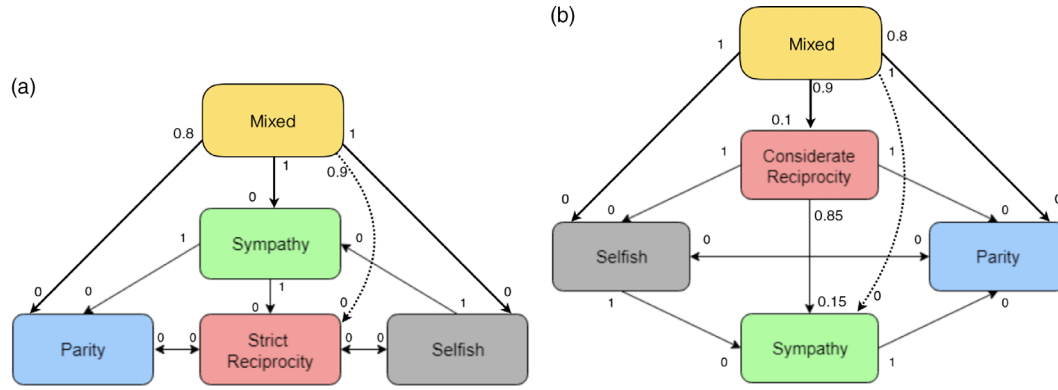


Figure 2 Prisoner's Dilemma: Dominance pattern in Head-to-Head interactions for (a) Strict Reciprocity and (b) Considerate Reciprocity. The numbers next to an agent type on an arrow represent the fraction of simulations where the average payoff of agents of that type that performed better than agents of the other type in the population (Note that the fractions on both ends of an arrow need not be 1 as the average payoff to both agent types may be identical in some runs.).

performed better than sympathetic and the latter in turn performed better than parity players (see Figure 2(a)). In the case where reciprocative players were considerate, sympathetic players performed rather poorly, only dominating the parity agents (causing a cyclic pattern with selfish, parity, and sympathy agents). Considerate reciprocative agents performed better than all but the Mixed agent in this case (see Figure 2(b)), an incredible improvement over the strict variant. These improvements result because like sympathy and mixed, considerate reciprocity converges to CC with itself, a trait that is not present in strict reciprocity. Yet considerate reciprocity still maintains the punitive aspect of strict reciprocity, by converging to DD with selfish, and DD with parity when parity plays first (opens with a D), CC when it plays second (reciprocity opens with a C).

4.2.1 Strict reciprocity

Reciprocity, parity, and selfish players perform the same because they all converge to Defect-Defect strategy profiles, which give 0 payoffs. This happens because all three types have $EU[D] > EU[C]$ as a first player and hence defects for the first move in the game. Also, as the second player, all three types defect as a response to the first player defecting. Defect-Defect is then the equilibrium strategy profile which results in a payoff of 0 for all parties involved.

Sympathy, on the other hand, always cooperates regardless of the other player's strategy. Therefore, it converges to a state of CC with parity, reciprocity and other sympathy players. All three of these groups cooperate as second players in response to a first player cooperating. The reason reciprocity and parity converge to this state instead of defecting with sympathy is that $EU[D] < EU[C]$ for both. Sympathy receives a higher payoff because it can additionally cooperate with other sympathy players and receive a higher payoff than what parity and reciprocity receive for interacting with their own types (0 payoffs from Defect-Defect outcomes).

4.2.2 Considerate reciprocity

The only difference here is that reciprocity performs better than parity and selfish players because considerate reciprocity players cooperate as the first player. In a group with parity players, the reciprocity players have CC outcome with themselves and CC with parity players, whereas parity players will have DD outcomes with themselves and with reciprocity players, thus receiving a lower payoff than reciprocity. When interacting with selfish players, reciprocity initially tries to cooperate, but chooses to defect when they find the selfish always defect. This way the only payoffs received apart from 0 payoff is when reciprocity interact with their own type and hence receive a higher payoff than selfish players. With sympathy players, the two types perform equally well because they CC with players of the same kind and the other kind.

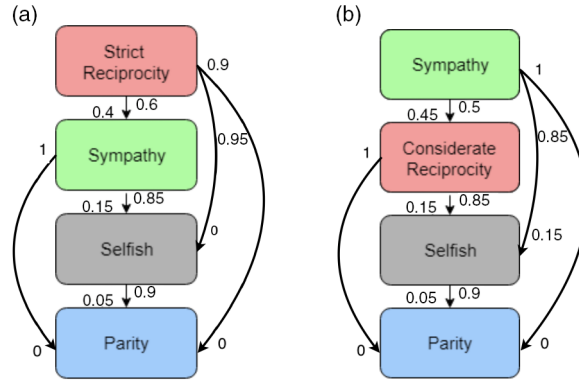


Figure 3 Battle of Sexes: Dominance pattern in Head-to-Head interactions for Strict Reciprocity (a) and Considerate Reciprocity (b)




Figure 4 Prisoner's Dilemma: Dominance pattern for Heterogeneous communities with Strict Reciprocity (a) and Considerate Reciprocity (b).

In terms of the Battle of Sexes game, for both Strict reciprocity and Considerate reciprocity, selfish and parity agents performed worst with selfish outperforming parity. The difference lies in the fact that sympathy underperformed to Strict reciprocity but outperformed Considerate reciprocity, as shown in Figure 3. It is clear to see that the head-to-head dominance patterns for the two games differ drastically, not only in terms of order but also in terms of the structure of the dominance pattern as there are no cyclic patterns seen in the Battle of Sexes game.

4.3 Heterogeneous communities

To study more realistic scenarios, we experiment with social networks where players of diverse types interact.

4.3.1 Heterogeneous communities (Prisoner's Dilemma scenario)

Dominance patterns in populations of sympathetic, strict reciprocating, parity and selfish agents:

When a social network contains equal distributions of each of these four groups, there is a clear order of dominance seen for the Prisoner's Dilemma game, as shown in Figure 4(a). The degree of each agent type reveals whom they interact with and why they gain such payoffs. For sake of conciseness, Figure 7 shows only the total degree of each agent type. It can be seen that Sympathy agents have the highest degree while Selfish the lowest and despite this Selfish agents perform better than Parity and Reciprocity. The reason behind this is, although all agents try to cut off all connections except those with Sympathy, the Selfish agents stand to gain a higher payoff in each interaction with Sympathy agents. Strict Reciprocative agents

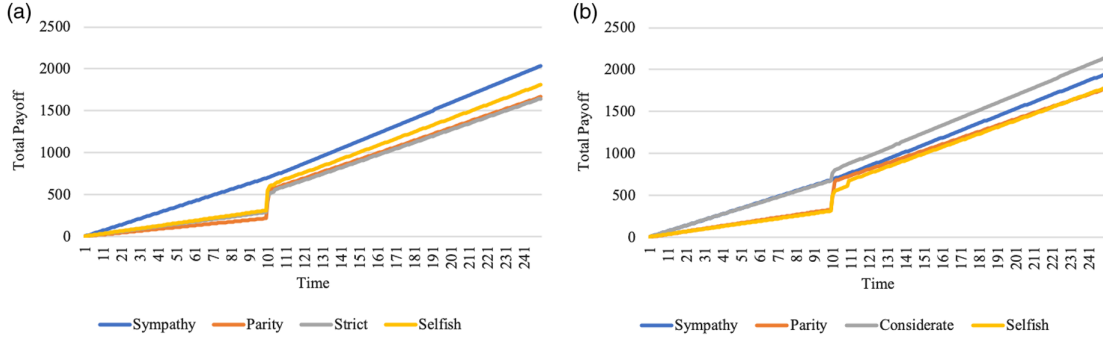


Figure 5 Prisoner's Dilemma Cumulative Payoffs: Strict Reciprocity (a) and Considerate Reciprocity (b).

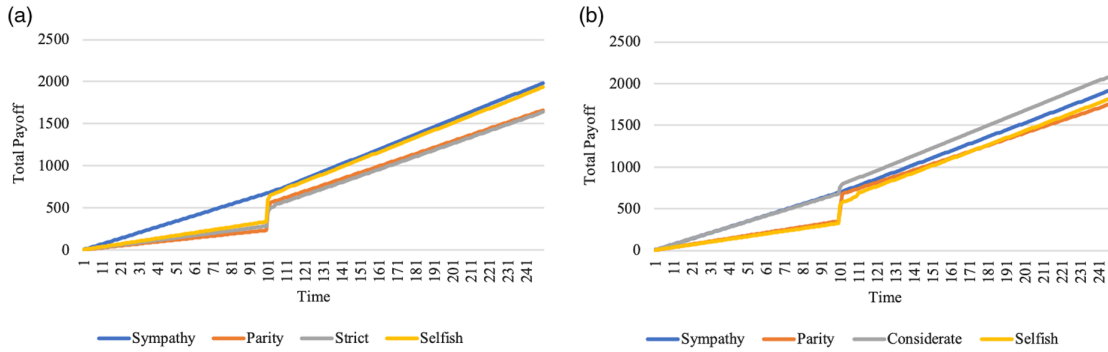


Figure 6 Prisoner's Dilemma cumulative payoffs: $\epsilon = 0.5$ and $\alpha = 5$ for Strict Reciprocity (a) and Considerate Reciprocity (b)

also gain higher than Parity agents in interactions with Sympathy agents but they lose out by remaining connected to some Reciprocative agents.

Dominance patterns in populations of sympathetic, considerate reciprocating, parity and selfish agents: The dominance pattern of these four groups, shown in Figure 4(b), is similar to the case with Strict Reciprocity except that Considerate Reciprocity now outperforms all other agent types. Such dominance can be attributed to the fact that Considerate Reciprocity players, unlike in the case of Strict Reciprocity, gain a high payoff from self-play. This, added to the payoffs received when interacting with Sympathy players, results in them gaining the highest payoffs compared to the other three agent types.

Cumulative payoffs: The graphs in Figure 5 show the average cumulative payoffs for each agent type over the rounds, averaged over ten trials, in a heterogeneous community. The dominance patterns observed in Figure 4 are evident in (a) Strict and (b) Considerate Reciprocity cases here. We observe that the Sympathy agents follow the same linear increase in their cumulative payoff, even when rewiring option becomes available after 100 time steps. All other agent types, however, receive a noticeable utility boost from re-wiring. We observe that agent types with high average cumulative payoffs also tend to have high degrees of connectivity in the network (see Figure 7).

We also present results from a second set of experiments in the Prisoner's Dilemma domain where ϵ value was increased to 0.5 from 0.1. This means the payoffs for player one would change to: $CC = 5$, $CD = -0.5$, $DC = 5.5$, and $DD = 0$. The dominance patterns remain unchanged in both cases of Strict and Considerate Reciprocity (Figures 6(a) and (b) respectively), despite the increased ϵ value.

4.3.2 Heterogeneous communities (Battle of the Sexes scenario)

The dominance patterns for cases of Strict and Considerate Reciprocity in the Battle of Sexes domain are presented in Figure 8. The corresponding cumulative payoff graphs for agents in heterogeneous groups interacting in the Battle of Sexes scenarios are shown in Figures 9.

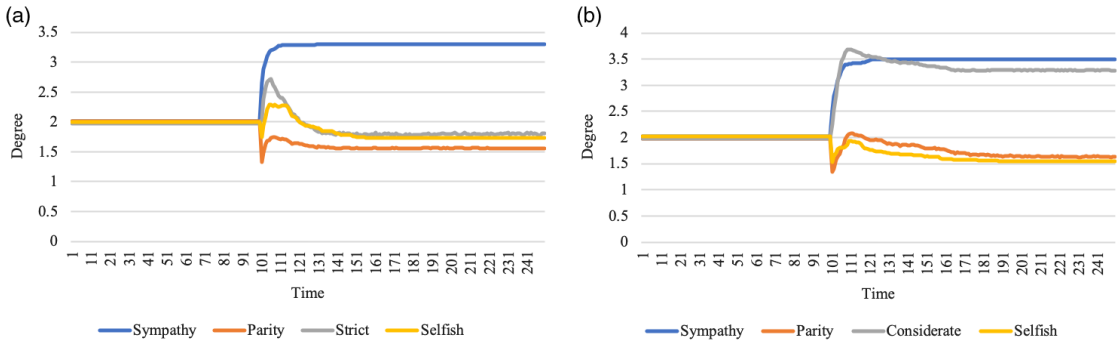


Figure 7 Prisoner's Dilemma: Average degree of agent type for the cases of Strict Reciprocity (a) and Considerate Reciprocity (b).




Figure 8 Battle of the Sexes: Dominance pattern for heterogeneous communities with (a) Strict Reciprocity and (b) Considerate Reciprocity.

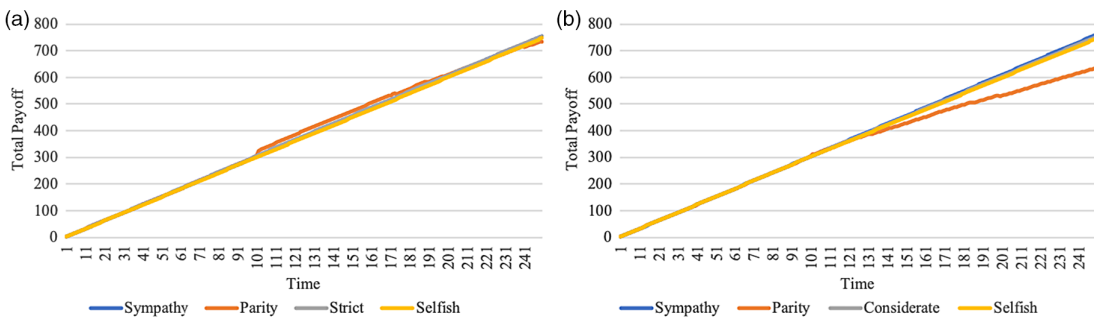


Figure 9 Battle of the Sexes cumulative payoffs: For Strict Reciprocity (a) and Considerate Reciprocity (b), averaged over ten trials.

The change in domain results in a partial change in the dominance pattern, where Sympathy, Selfish, and Parity agents retain their relative dominance positions in the Prisoner's Dilemma game (see Figure 4), across both cases of Reciprocity in the Battle of the Sexes game. The stark difference in dominance patterns in heterogeneous groups for Battle of the Sexes, when compared to that in the Prisoner's Dilemma scenario, is that here Strict Reciprocity performed best (see Figure 8(a)) as opposed to being the worst performing in Prisoner's Dilemma (see Figure 4), and the Considerate Reciprocity performed only second best (see Figure 8(b)) compared to being the best in the Prisoner's Dilemma scenario.




Figure 10 Topologies for increasing connection cost from (a) to (d): (a) Connection cost = -0.05 , (b) Connection cost = 1 , (c) Connection cost = 3.5 , (d) Connection cost = 5 .

4.4 Effect of connection cost on network topology

Connection cost, Γ has a significant impact on the topology of the network. Figure 10 shows the resultant topologies with heterogeneous agent populations for different connection costs in the Prisoner's Dilemma scenario. The green, red, blue, and black nodes correspond to Sympathy, Reciprocity, Parity, and Selfish agents respectively. A green edge represents a CC outcome, a red edge represents a DD outcome, a blue edge represents a DC outcome, and a black edge represents a CD outcome.

The topology becomes sparse with higher connection cost (see Figure 10). This is due to the fact that the players break off connections with opponent players when their $EU[C]$ or $EU[D]$ is lower than the connection cost. Increasing the connection cost reduces the incentive or tolerance for remaining connected to opponent players providing lower utilities.

5 Predicting average cumulative payoff for agent types

We now develop an analytical framework to predict average payoffs from interactions in environments containing our six agent types: selfish ($type_{me}$ or $type_1$), sympathetic ($type_s$ or $type_2$), parity ($type_p$ or $type_3$), strict reciprocity ($type_{sR}$ or $type_4$), considerate reciprocity ($type_{cR}$ or $type_5$), and mixed ($type_{mix}$ or $type_6$). The following analysis uses the Prisoner's Dilemma (PD) game; an identical analysis using the general formulas provided can also be performed for any other game, including the Battle of Sexes (BoS).

There are only two possible equilibria for each matchup, and such an equilibrium will be reached on the first or second interaction. This conclusion can be derived from a tree diagram of all possible outcomes of two interactions given the following observations for our six agent types in the PD: A decision to cooperate or defect is deterministic and depends only upon P_c , the fraction of cooperative responses by an agent's partner (and P_d , which is simply $1 - P_c$). If an agent type cooperated on *a priori* interaction with $P_c = p$, an agent type will always cooperate for any $P_c \geq p$. Similarly, if an agent type defected on *a priori* interaction when $P_c = p$, an agent type will always defect for any $P_c \leq p$. Equilibrium strategy outcomes for every combination of agent types are mapped below into 6×6 matrices (EP^i and EP^j) with 1–6 corresponding to agent types selfish, sympathetic, parity, strict reciprocity, considerate reciprocity, and mixed in that order. These outcomes were found by simply playing each matchup through with both orderings and recording the results.

The following four matrices are constructed using values achieved at the two possible equilibria. In the case of payoff at an equilibrium, any particular outcome (CC, CD, DC, DD) will always produce the same payoffs for every interaction, thus there is no dependence on the interaction number m . However, expected utility is a function of m , as agents calculate expected utilities based on P_c and P_d : values determined by outcome histories that differ between m values. In most cases for this particular analysis, both agents will cooperate every interaction or defect every interaction, which eliminates their dependence on m as P_c and P_d will be constants for all m . In the fewer cases where an equilibrium is achieved at interaction 2, expected utilities will depend upon m as P_c and P_d will not be constants for all m . We now introduce some terminology:

$EP_{i,j}^i$: Equilibrium payoff for type i when i leads in an interaction between types i and j .

$EP_{i,j}^j$: Equilibrium payoff for type j when i leads in an interaction between types i and j .

For $\alpha = 5$, $\epsilon = 0.1$, we get

$$EP^i = \begin{bmatrix} 0 & 5.1 & 0 & 0 & 0 & 0 \\ -0.1 & 5 & 5 & 5 & 5 & 5 \\ 0 & 5 & 0 & 0 & 0 & 0 \\ 0 & 5 & 0 & 0 & 0 & 0 \\ 0 & 5 & 5 & 5 & 5 & 5 \\ 0 & 5 & 5 & 5 & 5 & 5 \end{bmatrix}, \quad EP^j = \begin{bmatrix} 0 & -0.1 & 0 & 0 & 0 & 0 \\ 5.1 & 5 & 5 & 5 & 5 & 5 \\ 0 & 5 & 0 & 0 & 0 & 0 \\ 0 & 5 & 0 & 0 & 0 & 0 \\ 0 & 5 & 5 & 5 & 5 & 5 \\ 0 & 5 & 5 & 5 & 5 & 5 \end{bmatrix}$$

$EEU_{i,j}^i(m)$: Equilibrium expected utility for type i at the m th interaction between types i and j that occurs when i leads.

$$EEU^i(m) = \begin{bmatrix} 0 & 5.1 & 0 & 0 & 0 & 0 \\ 2.5 & 5 & 5 & 5 & 5 & 5 \\ 0 & 2.5 & 0 & 0 & 0 & 0 \\ 0 & 5 & 0 & 0 & 0 & 0 \\ 0 & 5 & 5 & 5 & 5 & 5 \\ 0 & 3.75 & 3.75 & 3.75 & 3.75 & 3.75 \end{bmatrix}$$

$EEU_{i,j}^j(m)$: Equilibrium expected utility for type j at the m th interaction between types i and j that occurs when i leads.

$$EEU^j(m) = \begin{bmatrix} 0 & 2.5 & 0 & 0 & 0 & 0 \\ 5.1 & 5 & 2.5 & 5 & 5 & 3.75 \\ 0 & 2.5 + \frac{2.5n - 2.5}{n} & 0 & 0 & 0 & 0 \\ 0 & 2.5 + \frac{2.5n - 2.5}{n} & 0 & 0 & 0 & 0 \\ \frac{5.1}{n} & 5 & 2.5 & 5 & 5 & 3.75 \\ \frac{5.1}{n} & 5 & 2.5 & 5 & 5 & 3.75 \end{bmatrix}$$

Given these matrices, we can construct two matrices describing the average interaction payoff for each type within every matchup prior to and during the rewiring phase.

Before the rewiring phase, an average interaction payoff can be described as a probabilistic weighting between the two possible equilibria payoffs. The expected interaction payoff without rewiring (EIP \bar{R}) for $type_i$ being matched with $type_j$ follows:

$$EIP\bar{R}_{i,j}^i = Pr(i)EP_{i,j}^i + Pr(j)EP_{j,i}^j,$$

where $Pr(i)$ is the probability that $type_i$ leads.

As it was determined that these two equilibria are determined by whether i or j leads, and because there is no bias in the topology generation process that will lead to some types having higher initial connections than other types, we can assume that each type is equally likely to initially contact the other, giving us $Pr(i) = Pr(j) = 0.5$. This gives us everything we need to construct our matrix of expected interaction payoffs without rewiring:

$$EIP\bar{R}^i = \begin{bmatrix} 0 & 5.1 & 0 & 0 & 0 & 0 \\ -0.1 & 5 & 5 & 5 & 5 & 5 \\ 0 & 5 & 0 & 0 & 2.5 & 2.5 \\ 0 & 5 & 0 & 0 & 2.5 & 2.5 \\ 0 & 5 & 2.5 & 2.5 & 5 & 5 \\ 0 & 5 & 2.5 & 2.5 & 5 & 5 \end{bmatrix}$$

The case after rewiring is a matter of *stable connections*, which are edges between agent types that both have expected utilities for the other agent that is higher than the connection cost $\Gamma = 0.5$. One or both of them would eventually disconnect otherwise. We make the simplifying and only partially correct assumption that unstable connections are immediately severed when the rewiring stage begins at round $R_{\bar{r}} + 1 = 101$. This assumption trades off a small degree of accuracy to avoid the great complexity of predicting the most likely sequence of disconnections each type will perform. We use an effective AND gate composed of two Heaviside (unit step) functions, H , that will produce 1 if and only if both agents will stay wired in a particular equilibrium, and 0 otherwise (indicating a disconnection) to compute the matrix for expected interaction payoffs after rewiring at interaction m with connection cost Γ :

$$EIPR_{i,j}^i(m, \Gamma) = Pr(i)EP_{i,j}^i H(EEU_{i,j}^i(m) - \Gamma) H(EEU_{i,j}^j(m) - \Gamma) \\ + Pr(j)EP_{j,i}^j H(EEU_{j,i}^j(m) - \Gamma) H(EEU_{j,i}^i(m) - \Gamma)$$

$$EIPR^i \left(m = 2 \frac{R_{\bar{r}}}{n}, \Gamma = 0.5 \right) = \begin{bmatrix} 0 & 5.1 & 0 & 0 & 0 & 0 \\ -.1 & 5 & 5 & 5 & 5 & 5 \\ 0 & 5 & 0 & 0 & 2.5 & 2.5 \\ 0 & 5 & 0 & 0 & 2.5 & 2.5 \\ 0 & 5 & 2.5 & 2.5 & 5 & 5 \\ 0 & 5 & 2.5 & 2.5 & 5 & 5 \end{bmatrix}.$$

Where interaction count m is assumed to be equal to the expected number of interactions performed at the time of rewiring decision: twice that of the number of pre-rewiring rounds divided by the mean degree n defined in the Watts-Strogatz construction. The factor of 2 comes from two opportunities for a connection in a single round to cause an interaction (one for each agent), and the factor of $\frac{1}{n}$ to give us the chance an opportunity actually causes an interaction (agents select an interaction partner uniformly between connections). When multiplied by our pre-rewiring rounds $R_{\bar{r}}$, we have the expected number of interactions prior to rewiring.

Our matrix after rewiring is coincidentally equal to the case without rewiring, although this could easily not be the case for different games, agent type definitions, connection costs, or pre-rewiring interaction counts.

These matrices will allow us to construct a formula for the average total payoff for every agent type. Our simplifying assumption is to consider an initially fully connected graph $K_{|\mathcal{A}|}$ of agents \mathcal{A} , whose total type payoff ratios will hopefully be representative of the mean total type payoff ratios resulting from all possible initial topologies. To arrive at expected total payoff for types, we will make use of the discovered expected interaction payoffs before and after rewiring and transform them into expected payoffs for rounds rather than interactions. We will begin with the post-rewiring case, as it turns out its result will inform us of the simpler case of pre-rewiring (where every agent is connected to every other agent). The expected payoff for an agent a of *type_i* in a round can be divided into two parts: the expected payoff from the one interaction initiated by a and the expected payoff from possible interactions with a initiated by other agents. The total stable connections (SC) an agent has will be necessary to know in order to calculate probabilities for type interactions. We can define the expected SC for an agent a of *type_i* as:

$$SC_i = \sum_{j=1}^{types} .5 \left(H(EEU_{i,j}^i - \Gamma) H(EEU_{i,j}^j - \Gamma) + H(EEU_{j,i}^i - \Gamma) H(EEU_{j,i}^j - \Gamma) \right) |type_j - \{a\}|$$

An expected payoff for the interaction initiated by agent a of *type_i* is a linear combination of probabilities of a *type_j* agent being sampled from a uniform selection of stable connections, multiplied by the expected interaction payoff for *type_i* being matched with *type_j*:

$$\sum_{j=1}^{types} \frac{|type_j - \{a\}|}{SC_i} EIPR_{i,j}^i(m, \Gamma) \quad (1)$$

The expected payoff for an agent a of $type_i$ received from interactions initiated by other agents is a linear combination of the amount of other agents of $type_j$, the chance for a to be sampled by a $type_j$ agent (SC_j^{-1}), and the expected interaction payoff for $type_i$ being matched with $type_j$:

$$\sum_{j=1}^{types} |type_j - \{a\}| \frac{1}{SC_j} EIPR_{i,j}^i(m, \Gamma) \quad (2)$$

Combining expressions (1) and (2), we find the expected round payoff for an agent a of $type_i$ after rewiring:

$$ERPR_i(m, \Gamma) = \sum_{j=1}^{types} \left(\frac{1}{SC_i} + \frac{1}{SC_j} \right) |type_j - \{a\}| EIPR_{i,j}^i(m, \Gamma).$$

For the pre-rewiring phase, we know that connection count is equal to $|\mathcal{A} - \{a\}|$ for an agent a regardless of type. This value will assume the roles of SC_i and SC_j which were previously necessary to consider disconnections. Additionally, $EIPR_{i,j}^i$ will be swapped for its pre-rewiring counterpart $EIPR_{i,j}^{\bar{i}}$:

$$\begin{aligned} ERPR_{\bar{i}} &= \sum_{j=1}^{types} \left(\frac{1}{|\mathcal{A} - \{a\}|} + \frac{1}{|\mathcal{A} - \{a\}|} \right) |type_j - \{a\}| EIPR_{i,j}^{\bar{i}} \\ &= \sum_{j=1}^{types} \frac{2|type_j - \{a\}| EIPR_{i,j}^{\bar{i}}}{|\mathcal{A} - \{a\}|} \end{aligned}$$

Expected total payoff for an agent a of $type_i$ is then equal to the expected total payoffs for each phase's rounds:

$$ETP_i = R_r * ERPR_{\bar{i}} + R_r * ERPR_i \left(m = 2 \frac{R_r}{n}, \Gamma \right).$$

Equal types: Considering a graph of 60 agents with 10 of each type we get the following expected payoff for different values of Γ (row i corresponds to $type_i$):

$$\Gamma = -1 \rightarrow \begin{bmatrix} 431 \\ 2069 \\ 845 \\ 845 \\ 1655 \\ 1655 \end{bmatrix}, \Gamma = 0.5 \rightarrow \begin{bmatrix} 1066 \\ 2683 \\ 1406 \\ 1406 \\ 2281 \\ 2281 \end{bmatrix}, \Gamma = 3 \rightarrow \begin{bmatrix} 171 \\ 2668 \\ 336 \\ 1499 \\ 2213 \\ 2213 \end{bmatrix}$$

Generalizations are difficult to claim with certainty as there exist a very large number of variables. Regardless, in proportions of agent types that are equal, sympathy seems to perform well for all Γ . Selfish agents have a sweet spot with $\Gamma = 0.5$ as they will not stay connected to each other, and will not be disconnected from sympathetic agents. When $\Gamma = -1$, selfish agents are punished by their own kind, and so is everyone else for staying connected to them and getting fruitless or exploitative outcomes, and when $\Gamma = 3$ no one tolerates selfish agents once they have a choice, and the selfish agent receive only what they can glean from the pre-rewiring stage. We now discuss payoffs to different agent types in some diverse population configurations:

Disproportionately sympathetic: For a graph of 1000 agents with 900 sympathetic and 20 of every other type we get the following expected payoff for different values of Γ :

$$\Gamma = -1 \rightarrow \begin{bmatrix} 2297 \\ 2249 \\ 2302 \\ 2302 \\ 2400 \\ 2400 \end{bmatrix}, \Gamma = 0.5 \rightarrow \begin{bmatrix} 2373 \\ 2453 \\ 2362 \\ 2362 \\ 2432 \\ 2432 \end{bmatrix}, \Gamma = 3 \rightarrow \begin{bmatrix} 919 \\ 2481 \\ 921 \\ 2391 \\ 2453 \\ 2453 \end{bmatrix}$$

Every agent type receives incredible payoffs in a highly sympathetic system. All agents seem nearly comparable until $\Gamma = 3$ where sympathy agents end their tolerance of selfish free-riding and parity agents fail to see the value in maintaining a connection with anyone (their construction is pessimistic, as the parity factor can only decrease expected utility).

Disproportionately selfish: Considering a graph of 1000 agents with 900 selfish and 20 of every other type we get the following expected payoff for different values of Γ :

$$\Gamma = -1 \rightarrow \begin{bmatrix} 51 \\ 203 \\ 100 \\ 100 \\ 198 \\ 198 \end{bmatrix}, \Gamma = 0.5 \rightarrow \begin{bmatrix} 801 \\ 611 \\ 995 \\ 995 \\ 1589 \\ 1589 \end{bmatrix}, \Gamma = 3 \rightarrow \begin{bmatrix} 20 \\ 1821 \\ 40 \\ 1197 \\ 1630 \\ 1630 \end{bmatrix}$$

Low payoffs are ubiquitous in a predominantly selfish population. Sympathy agents do not have their usual high ranking payoffs in $\Gamma = 0.5$, when they are exploited by the vast selfish majority. As the majority of the population is selfish, they can only exploit a small portion of the total population, which in turn prevents them from achieving higher relative payoffs.

6 Conclusion and future work

This work models utility-maximizing agents that adjust their perception of raw payoff based on three internal motivational factors: *sympathy*, *parity*, and *reciprocity*. The addition of these considerations to agents' utility calculations reflect aspects of human biases in real-world interactions.

When agents are interacting and forming new relationships on a social network, interesting network dynamics emerge. Head-to-head comparisons of individual agent types reveal a cyclic dominance pattern. Under many configurations a purely selfish strategy proves to be counter-productive, and can actually underperform to other agent types. Sympathetic, parity, and reciprocating agents all increase social and individual welfare to varying degrees under different configurations. Furthermore, agents that mix all three motivational factors with selfishness, generally have higher levels of individual and social welfare impact as compared to agent types with only a single personality trait.

The connectivity of a population is strongly dependent on the cost of forming and maintaining social network connections. Agent types can help improve the connectivity of a population, for example, sympathetic agents often connect agents of many different types when the connection cost is high.

These conclusions are seen to be valid across not only in the Prisoner's Dilemma game but also the Battle of Sexes game.

Equilibrium analysis of expected utilities and payoffs provide a way to predict final payoffs and stable connections among varying populations of agent types. Agent type success depends on a complicated web of factors. Connection cost causes drastic changes to expected agent type payoffs. Selfish populations cause drastic decreases in social welfare, while the opposite effect was demonstrated for sympathetic populations.

It will be interesting to observe how these conclusions will translate to sequential interaction scenarios, such as the *Investment Trust Game* (Berg *et al.* 1995; McCabe *et al.* 2001), when an agent's perception of their utility is influenced by sympathetic, reciprocative, or parity considerations.

Future work can include the development of formal predictions of emergent configurations given other system parameters such as initial network topology and stochastic interaction outcomes. We can

also study the dynamics of agent populations where agents can change their types based on observed performance on other agents.

Another fruitful research avenue would be to identify agent types that can be introduced into a population to reach desirable network configurations. Some examples of such ‘social network engineering’ include introducing agents into a selfish population to incentivize cooperative behavior or ‘toughening up’ a population of sympathetic agents to avoid manipulation from malicious selfish agents.

References

- Badhwar, N. 1993. Altruism versus self-interest: sometimes a false dichotomy. *Social Philosophy & Policy* **10**(1), 90–117.
- Baetz, O. 2015. Social activity and network formation. *Theoretical Economics* **10**(2), 315–340.
- Barabasi, A. 2016. *Network Science*.
- Berg, J., Dickhaut, J. & McCabe, K. 1995. Trust, reciprocity, and social history. *Games and Economic Behavior* **10**(1), 122–142.
- Bolton, G. 1991. A comparative model of bargaining: theory and evidence. *American Economic Review* **81**(5), 1096–1136.
- Bolton, G. E. & Ockenfels, A. 2000. ERC: a theory of equity, reciprocity, and competition. *American Economic Review* **90**(1), 166–193.
- Brooks, L., Iba, W. & Sen, S. 2011. Modeling the emergence and convergence of norms. In *IJCAI*, 97–102.
- Cha, M., Haddadi, H., Benevenuto, F. & Gummadi, K. P. 2010. Measuring user influence in twitter: the million follower fallacy. In *Proceedings of International AAAI Conference on Weblogs and Social, ICWSM '10*.
- Danassis, P. & Faltings, B. 2019. Courtesy as a means to coordinate. In *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems*. International Foundation for Autonomous Agents and Multiagent Systems, 665–673.
- David, E. & Jon, K. 2010. *Networks, Crowds, and Markets: Reasoning About a Highly Connected World*. Cambridge University Press.
- Dawes, R. M. & Thaler, R. H. 1988. Anomalies: cooperation. *Journal of Economic Perspectives* **2**(3), 187–197.
- Day, T. & Taylor, P. D. 1998. The evolution of temporal patterns of selfishness, altruism, and group cohesion. *The American Naturalist* **152**(1), 102–113.
- de Vos, H. & Zeggelink, E. 1994. The emergence of reciprocal altruism and group-living: an object-oriented simulation model of human social evolution. *Social Science Information* **33**(3), 493–517.
- Delgado, J. 2002. Emergence of social conventions in complex networks. *Artificial Intelligence* **141**(1–2), 171–185.
- Epstein, J. M. 2001. Learning to be thoughtless: social norms and individual computation. *Computational Economics* **18**(1), 9–24.
- Fehr, E. & Fischbacher, U. 2003. The nature of human altruism. *Nature* **425**(6960), 785.
- Fehr, E., Gächter, S. & Kirchsteiger, G. 1997. Reciprocity as a contract enforcement device: experimental evidence. *Econometrica: Journal of the Econometric Society* **65**(4), 833–860.
- Ferriere, R. & Michod, R. E. 1996. The evolution of cooperation in spatially heterogeneous population. *The American Naturalist* **147**(5), 692–718.
- Galán, J. M., Łatek, M. M. & Rizi, S. M. M. 2011. Axelrod’s metanorm games on networks. *PLOS ONE* **6**(5), 1–11.
- Gambetta, D. 1988. *Trust: Making and Breaking Cooperative Relations*. Blackwell.
- Goranson, R. & Berkowitz, L. 1966. Reciprocity and responsibility reactions to prior help. *Journal of Personality and Social Psychology* **3**(2), 227–232.
- Güth, W. & Yaari, M. E. 2004. Parity, sympathy and reciprocity. In *Advances in Understanding Strategic Behaviour*. Springer, 298–313.
- Krebs, D. 1970. Altruism – an examination of the concept and a review of the literature. *Psychological Bulletin* **73**(4), 258–302.
- Mahmoud, S., Miles, S. & Luck, M. 2016. Cooperation emergence under resource-constrained peer punishment. In *Proceedings of the 2016 International Conference on Autonomous Agents & Multiagent Systems, AAMAS '16*, Richland, S. C. (ed). International Foundation for Autonomous Agents and Multiagent Systems, 900–908.
- Martinez-Coll, J. & Hirshleifer, J. 1991. The limits of reciprocity. *Rationality and Society* **3**, 35–64.
- McCabe, K., Houser, D., Ryan, L., Smith, V. & Trouard, T. 2001. A functional imaging study of cooperation in two-person reciprocal exchange. *Proceedings of the National Academy of Sciences* **98**(20), 11832–11835.
- Peleteiro, A., Burguillo, J. C. & Chong, S. Y. 2014. Exploring indirect reciprocity in complex networks using coalitions and rewiring. In *Proceedings of the 2014 International Conference on Autonomous Agents and Multi-agent Systems, AAMAS '14*, Richland, S. C. (ed). International Foundation for Autonomous Agents and Multiagent Systems, 669–676.

- Rabin, M. 1993. Incorporating fairness into game theory and economics. *The American Economic Review* **83**(5), 1281–1302.
- Roth, A. E. 1995. Introduction to experimental economics'. In Kagel, A. and Roth, A. E. (eds). *Handbook of Experimental Economics*.
- Salehi-Abari, A. & Boutilier, C. 2014. Empathetic social choice on social networks. In *Proceedings of the 2014 International Conference on Autonomous Agents and Multi-Agent Systems, AAMAS '14*, Richland, S. C. (ed). International Foundation for Autonomous Agents and Multiagent Systems, 693–700.
- Sally, D. 1995. Conversation and cooperation in social dilemmas: a meta-analysis of experiments from 1958 to 1992. *Rationality and Society* **7**(1), 58–92.
- Schmitz, D. 1993. Reasons for altruism. *Social Philosophy & Policy*, **10**(1), 52–68.
- Sen, S. 1996. Reciprocity: a foundational principle for promoting cooperative behavior among self-interested agents. In *Proceedings of the Second International Conference on Multiagent Systems*, AAAI Press, 315–321.
- Sen, S. 2013. A comprehensive approach to trust management. In *Proceedings of the 2013 International Conference on Autonomous Agents and Multi-agent Systems, AAMAS '13*, Richland, S. C. (eds). International Foundation for Autonomous Agents and Multiagent Systems, 797–800.
- Taylor, P. 1992. Altruism in viscous populations – an inclusive fitness model. *Evolutionary Ecology* **6**, 352–356.
- Tsang, A. & Larson, K. 2014. Opinion dynamics of skeptical agents. In *Proceedings of the 2014 International Conference on Autonomous Agents and Multi-agent Systems, AAMAS '14*, Richland, S. C. (ed). International Foundation for Autonomous Agents and Multiagent Systems, 277–284.
- Watts, D. J. & Strogatz, S. H. 1998. Collective dynamics of 'small-world' networks. *Nature* **393**(6684), 440–442.
- Zhang, Y., Aziz-Alaoui, M., Bertelle, C. & Guan, J. 2014. Local nash equilibrium in social networks. *Scientific Reports* **4**, 6224.