

Effects of Parity, Sympathy and Reciprocity in Increasing Social Welfare

Rachna Nanda Kumar, Chad Crawford, Sandip Sen

The University of Tulsa

rachna-nandakumar@utulsa.edu, chad-crawford@utulsa.edu, sandip@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. 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 improving 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 maximizes 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 takes 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 and 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 and Ockenfels, 2000].

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

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 egoistical 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 investigate this perplexing question in the context of agents repeatedly interacting with neighboring agents located on a 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 recollect 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 that 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 net 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 unintuitive, 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) sympathy and parity are even more effective in improving social and individual welfare, (d) how influential ones 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 are not necessarily better off than 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.

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 [Barabasi, 2016; Baetz, 2015; David and Jon, 2010], such as 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 and Larson, 2014], or to characterize the emergence of conventions or norms [Delgado, 2002; Epstein, 2001] 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; Mahmoud *et al.*, 2016; Peleteiro *et al.*, 2014].

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 [Delgado, 2002; Epstein, 2001], e.g., in studies on emergence of conventions, the chosen stage games are symmetric coordination games with multiple Nash equilibria.

The behavioral traits of sympathy, parity and reciprocity has been inspired by a study involving human subjects [Güth and 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 takes into consideration the payoffs received by interaction partners, have been repeatedly observed to have key influence on human decision making. As a result researchers have found notable departures of human actors from self-seeking behaviors [Dawes and 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 and 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 and 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 [Badhwar, 1993; Day and Taylor, 1998; Krebs, 1970; Schmitz, 1993; Taylor, 1992]. Reciprocity [Goranson and Berkowitz, 1966], i.e., 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 [Fehr *et al.*, 1997; Martinez-Coll and Hirshleifer, 1991; Rabin, 1993].

Guth and Yaari [Güth and 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 and 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 opponent's past action, in promoting cooperation in repeated game scenarios and in multiagent societies [Ferriere and Michod, 1996; Sen, 1996; de Vos and Zeggelink, 1994]. 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 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 rewiring social network connections. The agents (players) are situated on a social network, where each agent

	C	D
C	(5, 5)	(-0.1, 5.1)
D	(5.1, -0.1)	(0, 0)

Figure 1: Raw Payoffs in a Prisoner’s Dilemma game.

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

Figure 2: Raw Payoffs in a Battle of Sexes game.

corresponds to a node and edges connect agents to other agents they can interact (play) with. The starting topology of the network is a *small world network*. [Watts and Strogatz, 1998]. A simulation run consists of repeated interactions of players over multiple iterations until a stable state is reached, i.e., no further changes in the network topology.

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 Figure 1, where $\alpha > 0$, $\epsilon > 0$ and $\alpha \gg \epsilon$. 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 Figure 2. The players move in turns: in this two-player, two-action game, the first player’s strategy choice is visible to the second player when it chooses its action.

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

$$u_w(x, y) = w_m x + w_s y - w_p |x - y| + w_r (\mathcal{C} \rho_c y - (1 - \mathcal{C}) \rho_d y)$$

where, $\mathcal{C} = 1$ when the other player cooperates and $\mathcal{C} = 0$ when the other player defects, ρ_c and ρ_d are the fraction of the other player’s payoff used by a reciprocative player when the second player cooperates and defects respectively.

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

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

where the four parameters represent the relative influences of sympathy, parity, reciprocity and selfishness, respectively, on the agent. The weights normalize the influences on utility, i.e., $w_s + w_p + w_r + w_{me} = 1$. We experimented with five types of agents; the corresponding weight vectors are presented in Table 1. The Utility matrix for each of these player types for each of the outcomes is presented in Table 2, as an example, for the Prisoner’s Dilemma game.

	$W_{sympathy}$	W_{parity}	$W_{reciprocity}$	W_{me}
Sympathy	0.5	0	0	0.5
Parity	0	0.5	0	0.5
Reciprocity	0	0	0.5	0.5
Selfish	0	0	0	1
Mixed	0.25	0.25	0.25	0.25

Table 1: Weight vectors of each of the four agent types.

	CC	CD	DC	DD
Sympathy	5	5.1	-0.1	0
Parity	5	-5.3	-0.1	0
Selfish	5	-0.1	5.1	0
Reciprocity $\rho_c = \rho_d$	10	-5.2	5	0
Reciprocity $\rho_c > \rho_d$	10	-2.65	5	0

Table 2: Utility to different player types for all possible game outcomes (first and letter in an outcome corresponds to the strategy choice of the first and second player respectively).

3.2 Types of players

Selfish: A player whose utility depends on only its own pay-off when deciding how to play with another player: $w_{me}=1$ and all other factors are assigned 0 weight.

Sympathy: A player whose utility is improved by an increase in the opponent’s payoff. We use $w_{me} + w_s=1$ and $w_r = w_p = 0$. For most experiments, $w_{me} = w_s = 0.5$, making Sympathy’s weight vector $\langle 0.5, 0, 0, 0.5 \rangle$.

Parity: A player that not only considers its own payoff but also tries to minimize the inequality in the payoffs between the two players: $w_{me} + w_p=1$ and $w_s = w_r = 0$. For most experiments, $w_{me} = w_p = 0.5$ and hence the Parity type’s weight vector is $\langle 0, 0.5, 0, 0.5 \rangle$.

Reciprocity: A player that includes reciprocative consideration of the opponent’s past strategy choice to determine its utility: $w_{me} + w_r=1$ and $w_s = w_p = 0$. For most experiments, $w_{me} = w_r = 0.5$ and hence the Reciprocity type’s weight vector is $\langle 0, 0, 0.5, 0.5 \rangle$. If the opponent cooperates (resp. defects), the reciprocative player adds (resp. subtracts) $w_r \rho_c$ (resp. $w_r \rho_d$), of the other player’s payoff to its payoff. We use two types of reciprocative players:

- **Strict Reciprocative** 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 Reciprocative** agents are less harsh to the opponent: given a uniform prior over opponent strategies, the player chooses cooperation over defection. Now

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

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

For a considerate reciprocative player, $EU[C] \geq EU[D]$, i.e., $\rho_c - \rho_d > \frac{2\epsilon}{\alpha + \epsilon}$. For $\alpha = 5$, $\epsilon = 0.1$, $\rho_c - \rho_d \geq 0.039$.

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’ pay offs and the other player’s strategy. Hence, the mixed player type is described by the weight vector $\langle 0.25, 0.25, 0.25, 0.25 \rangle$.

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, $P_d = t_d/(t_c + t_d)$.

The strategy chosen by a player, when making the first move, depends on its expected utility value of pay off for 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 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 and Strogatz, 1998] with parameters $n = 4$ and $p = 0.01$. 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 . 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, and hence forming/destroying connections will be the primary determinant of a user's influence in the network.

After 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, the agent a requests a recommendation from its "preferred partner," i.e., the neighbor it has the highest expected utility for, $EU_a = \max\{EU[C], EU[D]\}$, as the first player in the game. This "highest expected utility" is equivalent to a trust value for the agent: 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 with the highest expected utility from b that is not known to a . b reports the expected utility it has for c , EU_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 construct its own expected utility EU_c .

Following this recommendation stage, a elects to form a new connection with this neighbor if and only if the reported expected utility by b , 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 an "cognitive load" on the agent, such that interactions with others only pay dividends when the reward outweighs than the cognitive cost.

Disconnecting works as follows: after each interaction phase, each agent is given the choice to disconnect from their "least preferred" partner, d . If $EU_d \leq \Gamma$, the the agent will disconnect from this partner. Otherwise, the agent will keep all the connections they currently have.

4 Results

Unless otherwise stated, all games have been simulated with 100 agents for 250 iterations, allowing re-wiring after 100 iterations. Games were run for higher number of iterations and agents as well but lack of change in the results and better performance resulted in choosing these parameters as the standard.

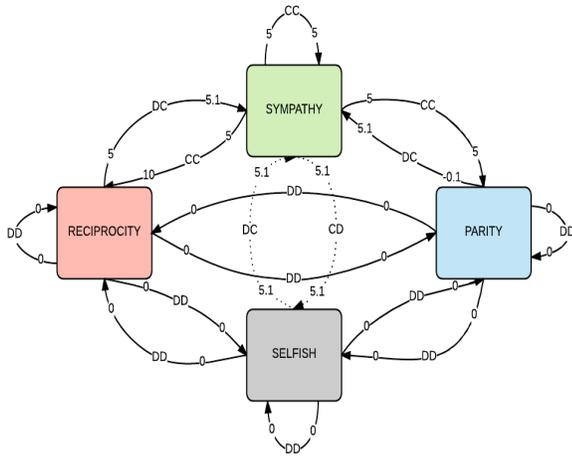
The results of the Prisoner's Dilemma game are presented in greater depth as the Battle of Sexes game was used primarily as a means to compare and contrast the results obtained with the Prisoner's Dilemma game to those in a different domain.

4.1 Pairwise interactions in one-on-one games

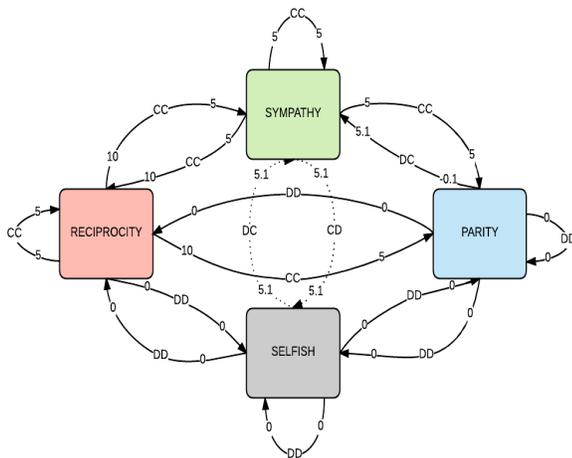
It is important to understand the convergence in players' strategies in repeated binary interactions, because those equilibrium strategy profiles will also be observed across the population when agents repeatedly interact with their neighbors on a social network. Convergence here refers to a Nash equilibrium, that is a state where each player has chosen a strategy that maximizes EU based on their observed distribution of action choices of the other player. The corresponding convergence strategies for Strict reciprocity and Considerate reciprocity are shown in Figure 3. 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, Reciprocity to Sympathy is CC instead of DC, and Reciprocity to Reciprocity is CC instead of DD.

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



(a) Strict Reciprocity case



(b) Considerate Reciprocity case

Figure 3: Equilibrium strategies with Strict Reciprocity

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 was a cyclic pattern in the dominance order of agent types in terms of who received the higher payoff in these head-to-head matches. Strict reciprocal players tied with parity and selfish players, all 3 types having an equal payoff when played against each other. However, parity and strict reciprocal agents performed worse than sympathy in pairwise interactions, while selfish agents performed better than sympathetic (see Figure 4 (top)). In the case where reciprocal players were considerate, sympathetic players performed rather poorly, only dominating the parity agents (causing a cyclic pattern with selfish, parity and sympathy agents). Considerate reciprocal agents performed best in this case, receiving a higher payoff than all the other types of agents (see Figure 4 (bottom)).

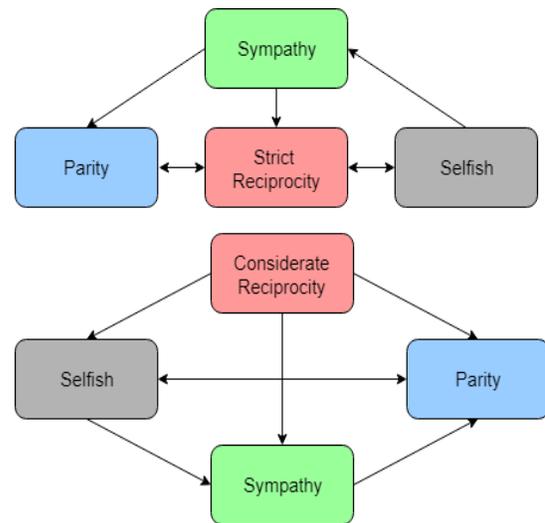


Figure 4: Prisoner's Dilemma: Dominance pattern in Head-to-Head interactions for Strict Reciprocity (top) and Considerate Reciprocity (bottom)

Strict reciprocity:

Reciprocity, parity and selfish players perform the same because they all converge to Defect-Defect strategy profiles, which gives 0 payoffs. This happens because all three types have $E[U(x)_d] > E[U(x)_c]$ 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 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 for both $U_{DC} < U_{CC}$. 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).

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.

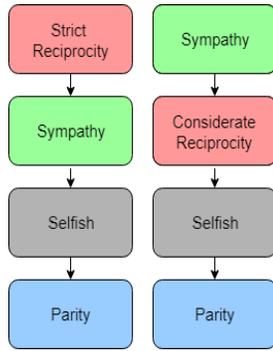


Figure 5: Battle of Sexes: Dominance pattern in Head-to-Head interactions for Strict Reciprocity (left) and Considerate Reciprocity (right)

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 5. 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 reflect a more realistic social network, where players of all types interact, we now evaluate social networks with all the agent types presented above.

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 6 (left). The degrees 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 degrees 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 also gain higher than Parity agents in interactions with Sympathy agents but they lose out by remaining connected to some Reciprocative agents.

Populations of sympathetic, considerate reciprocating, parity and selfish agents: The dominance pattern of these 4 groups can be seen in Figure 6 (right). It is similar to the case with Strict Reciprocity except that, here, Considerate Reciprocity outperforms Sympathy agents. Such dominance can be attributed to the fact that Considerate Reciprocity players here unlike in the case of Strict Reciprocity, gain a high

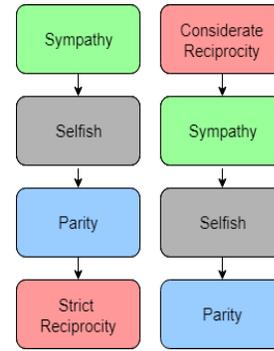


Figure 6: Prisoner's Dilemma: Dominance pattern for Heterogeneous communities with Strict Reciprocity (left) and Considerate Reciprocity (right)

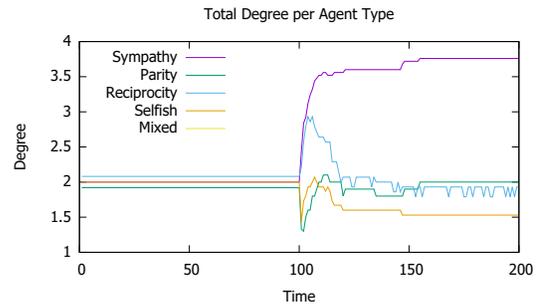


Figure 7: Prisoner's Dilemma: Average degree of agent type

payoff when they interact with themselves. This added to the payoffs received when interacting with Sympathy players results in them gaining the highest the payoffs compared to all other three types of agents.

The dominance patterns for both the above cases in the Battle of Sexes domain is presented in Figure 8. The change in domain results in a complete change in the dominance pattern, where Sympathy consistently performs the best and Parity performs the worst across both cases of Reciprocity in the Battle of Sexes game.

Mixed agents in heterogeneous communities

A heterogeneous community with mixed agents was simulated to test how their payoffs compared with that of other types of agents. In the case of strict reciprocative agents in the community, the dominance order observed before inserting mixed agents was changed: the selfish agents now receive the lowest pay off. In the case of considerate reciprocative agents, however, the dominance order that was observed before inserting mixed agents was maintained with the Mixed agents perform almost as well as the Reciprocity agents.

In the Battle of Sexes domain, the order of dominance with Strict and Considerate Reciprocity remained the same except that the Mixed agent performed the second best in both cases, constantly outperforming Parity, Selfish and Reciprocity agents while being dominated by Sympathy.

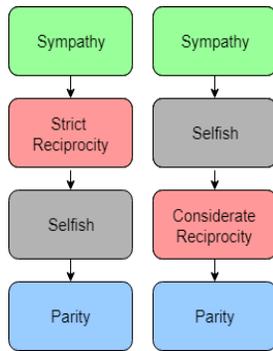


Figure 8: Battle of Sexes: Dominance Pattern for Heterogeneous communities with Strict Reciprocity (left) and Considerate Reciprocity (right)

4.4 Effect of connection cost on network topology

The connection cost has a significant impact on the topology of the network. Figure 9 shows the resultant topologies for different connection costs. 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 9). 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 their tolerance for remaining connected to opponent players providing lower utilities.

5 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, 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, e.g., 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

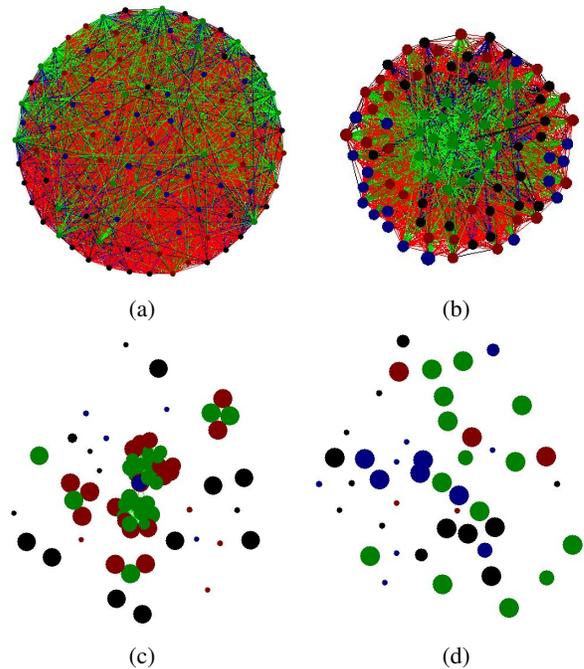


Figure 9: Topologies for increasing connection cost from (a) to (d): (a) Connection cost = -0.05, (b) Connection cost = 1, (c) Connection cost = 3.5, (d) Connection cost = 5.

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

Future work for this paper involves development of formal predictions of emergent configurations given initial population type distributions and other system parameters like network topology and connection cost. 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, 1993] N.K. Badhwar. Altruism versus self-interest: sometimes a false dichotomy. *Social Philosophy & Policy*, 10(1):90–117, 1993.
- [Baetz, 2015] Oliver Baetz. Social activity and network formation. *Theoretical Economics*, 10(2):315–340, 2015.
- [Barabasi, 2016] A.L. Barabasi. *Network Science*. 2016.
- [Berg *et al.*, 1995] Joyce Berg, John Dickhaut, and Kevin McCabe. Trust, reciprocity, and social history. *Games and Economic Behavior*, 10(1):122 – 142, 1995.

- [Bolton and Ockenfels, 2000] Gary E Bolton and Axel Ockenfels. ERC: A theory of equity, reciprocity, and competition. *American economic review*, pages 166–193, 2000.
- [Bolton, 1991] Gary Bolton. A comparative model of bargaining: Theory and evidence. *American Economic Review*, 81(5):1096–136, 1991.
- [Brooks *et al.*, 2011] Logan Brooks, Wayne Iba, and Sandip Sen. Modeling the emergence and convergence of norms. In *IJCAI*, pages 97–102, 2011.
- [Cha *et al.*, 2010] Meeyoung Cha, Hamed Haddadi, Fabrício Benevenuto, and Krishna P. Gummadi. Measuring user influence in twitter: The million follower fallacy. In *in ICWSM '10: Proceedings of international AAAI Conference on Weblogs and Social*, 2010.
- [David and Jon, 2010] Easley David and Kleinberg Jon. *Networks, Crowds, and Markets: Reasoning About a Highly Connected World*. Cambridge University Press, New York, NY, USA, 2010.
- [Dawes and Thaler, 1988] Robyn M. Dawes and Richard H. Thaler. Anomalies: Cooperation. *Journal of Economic Perspectives*, 2(3):187–197, September 1988.
- [Day and Taylor, 1998] Troy Day and Peter D. Taylor. The evolution of temporal patterns of selfishness, altruism, and group cohesion. *The American Naturalist*, 152(1):102–113, July 1998.
- [de Vos and Zeggelink, 1994] H. de Vos and E. Zeggelink. The emergence of reciprocal altruism and group-living: an object-oriented simulation model of human social evolution. *Social Science Information*, 33(3):493–517, 1994.
- [Delgado, 2002] Jordi Delgado. Emergence of social conventions in complex networks. *Artificial Intelligence*, 141(1–2):171 – 185, Jan 2002.
- [Epstein, 2001] Joshua M. Epstein. Learning to be thoughtless: Social norms and individual computation. *Computational Economics*, 18(1):9–24, 2001.
- [Fehr *et al.*, 1997] Ernst Fehr, Simon Gächter, and Georg Kirchsteiger. Reciprocity as a contract enforcement device: Experimental evidence. *Econometrica: journal of the Econometric Society*, pages 833–860, 1997.
- [Ferriere and Michod, 1996] Regis Ferriere and Richard E. Michod. The evolution of cooperation in spatially heterogeneous population. *The American Naturalist*, 147(5):692–718, 1996.
- [Galán *et al.*, 2011] José M. Galán, Maciej M. Łatek, and Seyed M. Mussavi Rizi. Axelrod’s metanorm games on networks. *PLOS ONE*, 6(5):1–11, 05 2011.
- [Goranson and Berkowitz, 1966] R.E. Goranson and L. Berkowitz. Reciprocity and responsibility reactions to prior help. *Journal of Personality and Social Psychology*, 3(2):227–232, 1966.
- [Güth and Yaari, 2004] Werner Güth and Menahem E Yaari. Parity, sympathy and reciprocity. In *Advances in Understanding Strategic Behaviour*, pages 298–313. Springer, 2004.
- [Krebs, 1970] D.L. Krebs. Altruism – an examination of the concept and a review of the literature. *Psychological Bulletin*, 73(4):258–302, 1970.
- [Mahmoud *et al.*, 2016] Samhar Mahmoud, Simon Miles, and Michael Luck. Cooperation emergence under resource-constrained peer punishment. In *Proceedings of the 2016 International Conference on Autonomous Agents & #38; Multiagent Systems*, AAMAS ’16, pages 900–908, Richland, SC, 2016. International Foundation for Autonomous Agents and Multiagent Systems.
- [Martinez-Coll and Hirshleifer, 1991] J. Martinez-Coll and J. Hirshleifer. The limits of reciprocity. *Rationality and Society*, 3:35–64, 1991.
- [McCabe *et al.*, 2001] Kevin McCabe, Daniel Houser, Lee Ryan, Vernon Smith, and Theodore Trouard. A functional imaging study of cooperation in two-person reciprocal exchange. *Proceedings of the National Academy of Sciences*, 98(20):11832–11835, 2001.
- [Peleteiro *et al.*, 2014] Ana Peleteiro, Juan C. Burguillo, and Siang Yew Chong. 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, pages 669–676, Richland, SC, 2014. International Foundation for Autonomous Agents and Multiagent Systems.
- [Rabin, 1993] Matthew Rabin. Incorporating fairness into game theory and economics. *The American Economic Review*, 83(5):1281–1302, 1993.
- [Roth, 1995] Alvin E Roth. Introduction to experimental economics’. in (a. kagel and ae roth, eds): *Handbook of experimental economics*, 1995.
- [Salehi-Abari and Boutilier, 2014] Amirali Salehi-Abari and Craig Boutilier. Empathetic social choice on social networks. In *Proceedings of the 2014 International Conference on Autonomous Agents and Multi-agent Systems*, AAMAS ’14, pages 693–700, Richland, SC, 2014. International Foundation for Autonomous Agents and Multiagent Systems.
- [Sally, 1995] David Sally. Conversation and cooperation in social dilemmas: a meta-analysis of experiments from 1958 to 1992. *Rationality and society*, 7(1):58–92, 1995.
- [Schmitz, 1993] D. Schmitz. Reasons for altruism. *Social Philosophy & Policy*, 10(1):52–68, 1993.
- [Sen, 1996] Sandip Sen. Reciprocity: a foundational principle for promoting cooperative behavior among self-interested agents. In *Proceedings of the Second International Conference on Multiagent Systems*, pages 315–321, Menlo Park, CA, 1996. AAAI Press.
- [Taylor, 1992] P.D. Taylor. Altruism in viscous populations – an inclusive fitness model. *Evolutionary Ecology*, 6:352–356, 1992.
- [Tsang and Larson, 2014] Alan Tsang and Kate Larson. Opinion dynamics of skeptical agents. In *Proceedings of the 2014 International Conference on Autonomous Agents and Multi-agent Systems*, AAMAS ’14, pages 277–284,

Richland, SC, 2014. International Foundation for Autonomous Agents and Multiagent Systems.

[Watts and Strogatz, 1998] Duncan J Watts and Steven H Strogatz. Collective dynamics of ‘small-world’ networks. *nature*, 393(6684):440–442, 1998.