

Behavioral Conflict and Fairness in Social Networks

Stephen Judd¹, Michael Kearns¹, and Yevgeniy Vorobeychik²

¹ Computer and Information Science, University of Pennsylvania, Philadelphia, PA

² Sandia National Laboratories, Livermore, CA

Abstract. We report on a series of behavioral experiments in social networks in which human subjects continuously choose to play either a dominant role (called a *King*) or a submissive one (called a *Pawn*). Kings receive a higher payoff rate, but only if all their network neighbors are Pawns, and thus the maximum social welfare states correspond to *maximum independent sets*. We document that *fairness* is of vital importance in driving interactions between players. First, we find that payoff disparities between network neighbors gives rise to conflict, and the specifics depend on the network topology. However, allowing Kings to offer “tips” or side payments to their neighbors substantially reduces conflict, and consistently increases social welfare. Finally, we observe that tip reductions lead to increased conflict. We describe these and a broad set of related findings.

1 Introduction

Reporting on a series of behavioral experiments involving a particular class of coordination tasks on social networks, we demonstrate the central importance of *fairness* and *conflict* in interactions between players which entail exclusively financial consequences. The experiments were held in a single session with 36 human subjects, each controlling the state of a single node in an exogenously imposed social network. In our first set of experiments, each subject could choose to be a *King* or a *Pawn*. A King is paid at a higher rate (twice as much as a Pawn), but only if no network neighbor is in *conflict* with him or her by also having chosen to be a King; a King in conflict receives no payments. Players can asynchronously change their state at any time. Since only one of any two neighbors can be a King for either to be paid, such a configuration is inherently “unfair”, giving rise to considerable tensions between pure self-interest and fairness considerations. Our second set of experiments thus involved an additional element: Kings that had no conflicts were able to designate a *tip* (or side payment) which was equally divided among their Pawn neighbors.

Networked King-Pawn games may broadly be seen as modeling economic interactions in which each local neighborhood can support only one dominant player. For example, in organized crime it is often the case that only one clan or faction is permitted to rule a locality, and incursions against the incumbent often result in violent clashes that are damaging to both sides. We may also

consider geographic sovereignty as an example — governments oversee property, and attempts by neighboring nations to overtake that property may result in costly wars.

A game theoretic understanding of dynamic coordination games such as these offers several approaches. One considers a stylized one-shot game modeling a long-run outcome, which in our setting exhibits no conflict in a pure strategy Nash equilibrium, and in which any positive tip level is strictly dominated. Another is a repeated game model, with a concomitant explosion of equilibria. Our behavioral results contradict the predictions of a one-shot model: tipping is marked when allowed, and persists even at the end of games. A repeated game model, on the other hand, has too many equilibria to offer a meaningful prediction, and does not suggest any fundamental difference between our two settings. In our experiments, on the other hand, we document numerous qualitative and quantitative differences between the first (no tips) and the second (with tips) settings.

One of our most notable observations is that social welfare is uniformly higher when tips are allowed. Even if we allow that such outcomes are consistent with *some* equilibrium of a repeated-game model, we must still appreciate why this, and not another, equilibrium is ultimately chosen by human subjects. We argue that (the lack of) perceived fairness is primarily responsible for the observed differences between the two experimental settings: inequality in payoffs creates considerable conflict in the first setting, and tips ameliorate conflict by bridging the payoff gaps in the second setting. This finding is robust to the network topology and has broad implications, including the suggestion that reducing income inequality may actually raise social welfare. While there has been well-documented evidence of the importance of fairness considerations in ultimatum games [7, 4], we note that our setting involves global coordination on networks, rather than bargaining, and it is not a priori clear that equity plays a role in facilitating or hampering coordination.

A well-known theory in macroeconomics suggests that wages are resistant to reduction, since people view reductions in their wages as inherently unfair even if their real value is preserved [1]. Roughly mapping tips in our setting to wages, we find behavioral evidence for this “downward rigidity”. Specifically, we observe that for similar average tip levels, a tip reduction resulted in considerably more conflict. Furthermore, we find that the amount of conflict in response to tip reductions actually rises with average tip pay rate—higher earners appear to respond more strongly to pay cuts.

The experiments described here are part of a broader and ongoing program of behavioral experiments in strategic and economic interaction on social networks conducted at Penn [14, 11, 13], and are an effort to apply the methods of behavioral game theory [6] to the study of social networks.

1.1 Related Literature

The games we study are networked generalizations of repeated or continuous versions of the game of *Chicken* or *Hawk-Dove* [9], 2-player instances and certain

generalizations of which have been studied extensively in the lab [15, 17, 3]. The subject of fairness in human interactions has a very long history as well. Sociologists and social psychologists view it as central to many social phenomena, and have well-developed theories of fair exchange and reciprocity (exchange/equity theory) [5]. The economic experiments of Fehr and Gächter [8] show that people frequently punish non-altruistic behavior and derive pleasure from doing so. Akerlof and Yellen develop a hypothesis of wage effort based on fairness considerations [2] which allows them to offer an explanation of unemployment and supports the general observation that wages tend to be downwardly rigid [1]. Rabin [18], Fehr and Schmidt [7], and Bolton and Ockenfels [4] offer alternative theories that incorporate fairness into more traditional game theoretic models.

The term “social welfare” will be used here to mean the total payoff to all players in a game. It is worth noting that maximizing the social welfare of our game is isomorphic to the Maximum Independent Set problem, which is a canonical NP-Complete problem [10]. In this study, we construct games in such a way that a Pareto optimal pure strategy Nash equilibrium of its one-shot version solves the maximum independent set problem. In that regard, this work is similar to the experiments in which subjects were placed as nodes in a graph and tasked with coordinating on a *proper coloring*—another canonical NP-Complete problem [14, 12].

2 Experimental Design

In our experiments players were mapped to nodes on exogenously specified networks. Each player was given one of two action (role) choices: to be a King or a Pawn. As a King, the player would enjoy a high pay rate (\$1/minute), but payments only accrue if *none of his neighbors are also Kings*. A *conflict* is a situation in which there are two neighbors both selecting King — and both earn zero. Being a Pawn, in contrast, is risk-free: no matter what their neighbors choose, Pawns earn a steady income, albeit only half of a King’s (\$0.50/minute). Payments accrued continuously for each player, pro-rated by the time spent in each of the three possible local states (King without conflicts, King with conflicts, and Pawn). Players could asynchronously update their choices at any time.

A conflict-free configuration of Kings forms an *independent set*. Since Kings are paid only when all their neighbors are Pawns, social welfare is maximized when Kings form a *maximum independent set*, though computing such a maximum is NP-Hard in general.

We ran two variants of this basic King-Pawn scenario. The first was precisely as described above. In the second, we allowed players to offer tips to each other. Tips are payable only while a King is *non-conflicting* (i.e., he is a King and all of his neighbors are Pawns), and when payable they are divided equally among neighbors. Tip offer values were an amount between 0 and 100% of a King’s pay rate, but were restricted to quantum steps of 10% (i.e., 10 cents/min). We call this second scenario the “tips” setting, in contrast to the former, which we call the “no-tips” setting.

A natural question to ask is whether allowing players to exchange tips is at all consequential according to traditional game theory. Let us thus observe that *in the tips setting, non-conflicting Kings should never offer tips at Nash equilibrium of a one-shot game corresponding to our setup*. This observation also holds in the last stage of finite repeated games. Since the experiments involve a known time limit and our clock has a finite granularity, we can view them as finite-period repeated games; in such a repeated game, positive tipping could indeed occur even in a subgame perfect equilibrium (except in the last stage), since a mixed strategy equilibrium of a stage game can offer a credible threat.

All experiments were held in a single session lasting multiple hours with 36 University of Pennsylvania students as participants. We ran two sets of 19 experiments, one set for the no-tips and another for the tips setting. Each experiment had a fixed network topology, and subjects were randomly assigned to nodes. All experiments lasted exactly two minutes.

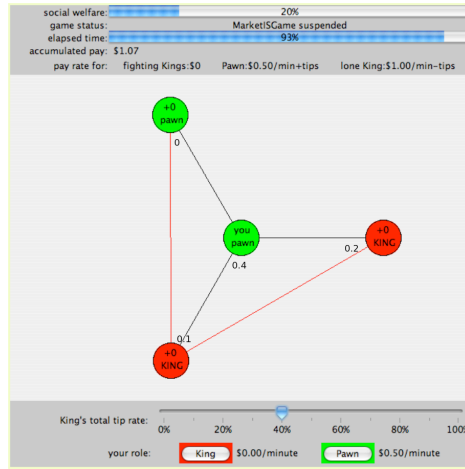


Fig. 1. A screenshot of a player’s GUI for the tips scenario. The central node represents the player using the GUI. The numbers displayed near the circles indicate tip offers. The slider designates a choice for the tip offer. The buttons at the bottom of the screen allow a player to choose to be a king or a pawn. In the no-tips setting all allusions to tips (including the slide bar and tip amounts near the nodes) are removed.

A screenshot of the tips GUI is shown in Figure 1; for the no-tips setting, the tip rate bar was simply absent. Each player could see his neighbors and relationships between them, as well as their role and tip choices, but could not see relationships or actions of anyone else. All actions were asynchronous. Role changes or tip adjustments could be made at any time during the game. The session was closely proctored and physical partitions were erected to ensure no communication between subjects.

In both the no-tips and the tips settings we ran three experiments each on six network topologies (Bipartite, Preferential Attachment Tree, Dense Prefer-

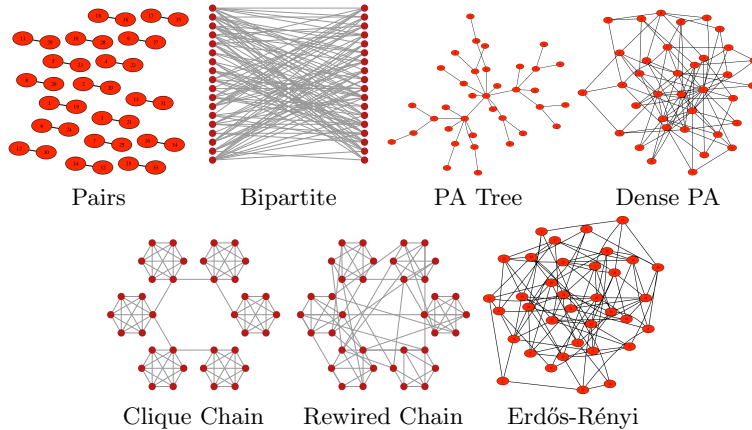


Fig. 2. Sample network topologies used in experiments.

ential Attachment, Clique Chain, Rewired Chain, and Erdős-Rényi) and a single experiment on a Pairs topology. Visualizations of typical candidates from each topology class are provided in Figure 2. If a specific topology is a class with a stochastic generative model (i.e., one of Bipartite, Preferential Attachment Tree or Dense Graph, Rewired Chain, and Erdős-Rényi Graph), we generated a different network in each of a set of three experiments on that topology, but used the same graphs in both the no-tips and the tips settings. In the Preferential Attachment (PA) Tree, each node that is added to the graph is connected to exactly one existing node. In the Dense PA topology, a new node is connected to three existing nodes. For Erdős-Rényi graphs, we set the probability p of an edge between two nodes to 0.15. Each of 102 edges in the Bipartite graphs paired players uniformly at random. Rewired Chain starts with a Clique Chain as a baseline and rerouts each intra-clique edge with probability 0.2. More detailed descriptions and motivation for these and similar generative models can be found in [16].

The *clustering coefficient* of networks is relevant to our results. It is defined as the number of closed triplets divided by the number of connected triplets of vertices. Figure 3 (left) and later bar plots organize the networks in increasing value of their clustering coefficient. In figures throughout, we mark a network by *** if the reported result or difference is significant with $P < 0.01$, while ** indicates $P < 0.05$, and * corresponds to $P < 0.1$ significance level. When such a result is attributed to both the no-tips and tips settings, we marked the pair with the lowest significance level observed.

3 Results

3.1 Collective Wealth and Tipping

Game-theoretic solutions (applied most directly) do not predict a fundamental difference arising from allowing players to exchange tips. Figures 3 and 4, how-

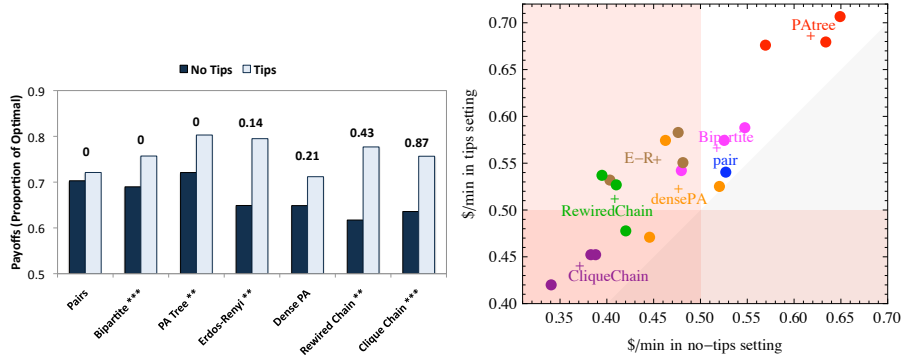


Fig. 3. Left: Achieved welfare (average payoff (proportion of optimum) per player per minute) in the no-tips (black) and tips (white) settings; networks ordered from left to right by clustering coefficient, which is also displayed above each set of bars. Right: Average rate of earnings in each game. Each of the 19 networks is shown without averaging into its replication group; all 19 of them fall above the gray triangle, indicating uniform improvement in the tips setting. The + marks are located at the averages of the replication groups. The shaded zones are where performance is below Pawn rate.

ever, demonstrate a systematic improvement in welfare under the tips setting. The impact of tips on welfare varies greatly, and is substantial for some networks.

In Figure 3 (left) we report the *relative* social efficiencies (behaviorally realized social welfare as a proportion of the theoretically optimal social welfare) for the different network topologies (averaged over trials), under both the no-tips and tips settings. Clique Chain, Rewired Chain, and Erdős-Rényi networks exhibit the greatest payoff improvements under tips (around 15% of optimal welfare). Note that these are also three of the four most clustered networks; more on that later. The Pairs network — where there is no “network” per se but rather 18 separate one-on-one games — shows the least improvement, suggesting that the social welfare benefits of tips increase with network complexity. Payoff improvements were significant ($P < 0.05$) in 5 of 7 network architectures (shown in Figure 3, left), and overall improvement in welfare was significant with $P < 0.01$. One may suggest that the reason for the improved outcomes in the tips setting can be attributed to learning effects. To rule out this explanation, we correlated the experiment sequence index with corresponding welfare outcome separately in the no-tips and tips settings. The correlation coefficient was small in the no-tips setting, somewhat larger when tips were allowed, but not statistically significant in either case; it seems clear in any case that subjects had not learned to play the game any better during the no-tips sequence.

Figure 3 (right) illustrates the *absolute* average rates of income for all 19 networks in each of the two settings. The PA trees stand out as being particularly wealthy in both settings; the CliqueChains performed below Pawn level in both settings. The ER graphs are all in the upper left quadrant; they all move from sub-Pawn losers to relative winners when tips are allowed. This figure also

demonstrates that not only the averages, but all 19 individual network topologies yielded higher social welfare under tips.

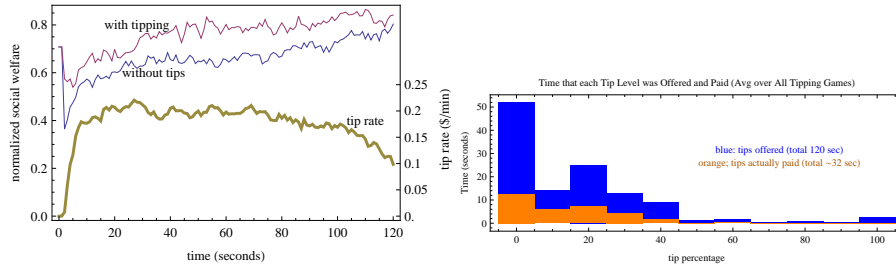


Fig. 4. Left: The top two lines are the average pay rate (as a proportion of optimum) over time in the two settings, averaged over all experiments. Their scale is on the left. The bottom (thick) line is the tip rate, averaged over all tips experiments. Its scale is on the right. Right: Distribution of tips, both offered and accepted.

Some of the aggregate dynamics (Figure 4, left) reveal the effects of tipping. Average welfare improvement due to tips is persistent throughout the span of experiment. Furthermore, we find that the use of tips, when allowed, is rather substantial: tips accounted for 13% of all income in the tips setting. The lowest curve in the figure shows the average tip rate offered by players over time (average taken over all tips experiments). The tip rate is initialized to zero, but jumps almost immediately after an experiment starts, and persists at around 20% for the bulk of the experiment. It falls off gradually between 70 and 100 seconds and then faster after the 100 second mark, but, even at the end of the experiments, average tip rate persists at around 10%, well above equilibrium level.

We observe that welfare rises over the span of a game in both settings. Furthermore, social welfare *increases* over the last 90 seconds, even as tipping *decreases*. This observation suggests an alternative hypothesis that tips serve as a coordination device, similar to cheap talk, to help select an equilibrium. While such an explanation seems difficult since it would require players to first coordinate on a global meaning of tips in order to use it as a communication device, and is further undermined by the observed persistence of non-negligible tipping at the ends of games, we cannot fully rule it out given our experimental design. We found no significant correlation between tip rates and experiment index within the session, suggesting no long-term adaptation of tipping behavior.

Figure 4 (right) shows the amount of time tip sliders spent in each of the 11 possible states (averaged over all players and games). In the case of one King adjacent to a single Pawn, the tip amount that divides the income equally is 25%. One of the modes in this histogram is slightly below that, at 20%. In the vaguely similar setting of the Ultimatum Game [6] there is a mode of 30 or 40%, also slightly below equitable.

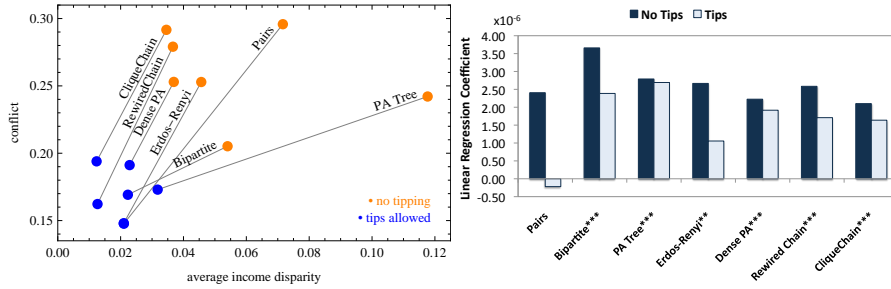


Fig. 5. Left: The connection between income disparity and conflict. Both quantities are reduced when tips are allowed. Right: The connection (linear regression coefficient) between income disparity and conflict (after normalizing both quantities) by network class for the no-tips and tips settings.

3.2 Conflict and Fairness

Thus far we have established the substantial use of tips when available, and their consistent improvement of social welfare. But what behavioral processes underly these phenomena? Here we propose and support the following hypothesis: *Subjects used conflict — which reduces the wealth of all players involved — to express perceived unfairness or inequality. Tipping reduces unfairness and consequently reduces conflict, thereby raising the average payoffs of all players and facilitating coordination.*

We begin this analysis by contrasting quantitative measures of income inequality between the no-tips and tips scenarios. Consider first just the horizontal axis of Figure 5 (left), which measures average income disparity (defined as the average squared difference in payoffs between network neighbors). Since tip levels persist well above zero, and that money is being routed to other players, it is significant and unsurprising that income disparity falls when tipping is allowed. What is more interesting is that tipping appears to roughly equalize payoff asymmetries *across networks*, which were substantially more variable in the no-tips case. For example, PA Tree networks that had shown large income inequality in the no-tips setting are now much closer to other network types. We found a significant correlation (0.49, $P < 0.04$) between income disparity under the no-tips setting and tips exchanged when they are allowed. Our interpretation is that the more a game is perceived as unfair, the greater the role that tips must play in bridging income gaps between players.

The role of tipping in reducing income inequality is only one part of our hypothesis. Additionally, we posit that conflict expresses a *perception* of unfairness. Since tips reduce inequity, we propose that they alleviate the tension that leads to conflict; thus, tips effectively replace or substitute for conflict when they bridge inequality gaps. To support the idea that tips substitute for conflict, we expect to see substantial reduction in conflict between players when tips are allowed. Figure 5 (left) shows this on its vertical axis: the amount of conflict between players (specifically, the average proportion of the game that a player spent in conflict, with average taken over players and games) is systematically

lower in the tips setting. Nevertheless, it is difficult to establish a clear relationship between income disparity and conflict. We conjecture that what matters is *perceived*, rather than *observed* (or measured) unfairness, as suggested by *equity theory* [5]. For example, it may seem fair that low degree nodes receive higher income due to the natural advantage their network position offers. We can test this conjecture by considering the correlation between income inequality or conflict with average disparity of degrees between network neighbors; however, we did not find such correlations to be significant in our setting. Instead, we found that the *clustering coefficient* exhibited significant correlation with time players spent in conflict in the no-tips setting (0.62, P -value < 0.01); correlation between the same quantities is considerably smaller and not significant in the tips setting.

As more direct support that conflict communicates perceived unfairness, we looked at *individual* level correlations between the time that a player spends in conflict that he initiates and ultimately terminates, and that player’s perceived income disparity, defined as zero when his income is higher than a neighbor’s and the squared payoff difference otherwise, and averaged over all of his neighbors.³ The correlation between these quantities is 0.345 ($P < 0.001$) in the no-tips setting and 0.25 ($P < 0.001$) in the tips setting. These correlations suggest that when players perceive unfairness in their predicament, they are much more likely to engage in conflict with neighbors. On the other hand, the correlation is markedly weaker in the tips setting, providing further evidence for substitution between conflict and tips. One may hypothesize that conflict serves the purpose of punishment to motivate coordinated, better outcomes, similar to Prisoner’s Dilemma; below we refute this by showing that conflict decidedly does not pay.

We next consider again the correlation between perceived income inequality and conflict, separated by individual network. However, rather than simply looking at correlations between the two quantities, we regress time a player spends in conflict on his perceived income disparity. In Figure 5 (right) we report the regression coefficient. The figure does not appear to exhibit much systematic difference in the linear relationship between conflict and perceived income disparity across networks. While there does appear to be a slight negative trend as the clustering coefficient increases in the no-tips setting, we did not find it to be statistically significant. Nevertheless, the relationship is clearly positive—significantly so in all graph classes except “Pairs”.

Conflict appears to also serve as a means of tip bargaining. Let C be the time (in seconds) that a player spends in conflict that he both initiates and terminates. Define T as *tip income rate*, that is, average tip income per minute that a player spends as a Pawn. Let W be the wealth of a player for the entire game. The correlation between C and T is 0.19 ($P < 0.001$), while the correlation between C and W is -0.51 ($P < 0.001$). The positive correlation between C and T generalizes across 5 of 7 network architectures (significant in all 5); the only exceptions are Clique Chain and Rewired Chain. Thus, while conflict may show

³ This definition of fairness closely mirrors the notion introduced by Fehr and Schmidt [7].

some success in negotiating a higher tip income rate, it yields an unambiguous loss in the long run.

To quantify the tradeoff between time spent in conflict and tip income rate, as well as conflict and wealth, we fit linear regression models to both sets of data pairs. We find (with coefficients having $P < 0.001$) that every second that a player engages his neighbors in conflict earns him (on average) an additional 0.2 cents in tips. Regressing wealth against conflict, on the other hand, tells us (with even higher significance for both regression coefficients) that every second in conflict *costs* a player 1.2 cents on average. The struggle for a bigger tip yields meager rewards and ultimately costs a player more than it is worth.

3.3 Downward Rigidity of Tips

One explanation of high unemployment offered in macroeconomic theory posits that wages are *downwardly rigid*, as people view wage decreases as unfair, even if these decreases maintain the real value of wages (e.g., when there is deflation) [2, 1]. As a result, employers prefer to offer above-market wages to ensure that worker productivity remains high; what results is a shortage of jobs relative to the number of people seeking work.

There is a suggestion in Figure 4 (left) that tip changes are downwardly rigid. After being quickly established at the 20% level, they are very slow to head toward equilibrium level and never fall even half way back to zero.

Figure 6 (left) supports the hypothesis that downward changes are viewed as unfair more directly. The comparisons in the figure are between players who made at least one tip reduction and those who made none. The players who did make a tip reduction suffered more conflict than those who did not, even as average tip income rates were roughly equal between the groups. Additionally, *as tip rates increase, tip reductions actually entail more, not less, conflict*. To test the significance of this, we looked at finer discretized tip income rate intervals and correlated midpoints of these with average increases in conflict time. The resulting correlation was 0.99 and highly significant ($P < 0.001$). This result cannot be explained by suggesting that higher tippers also made greater tip reductions: we found no significant correlation between tip pay rate and average size of a tip cut.

3.4 Individual Nodes

The previous discussion pertains to the communal patterns of behavior, but there were also interesting variations at the level of individual nodes.

One natural question to ask is whether a node's degree had an impact on its wealth and role choices. We found significant negative correlation between a node's degree and wealth in both settings (correlation of -0.33 in the no-tips setting, -0.26 in the tips setting, both with $P < 0.001$). Thus, having a high degree was, overall, a handicap. However, breaking this down by network class (Figure 6, right) we find that the negative relationship between degree and income is only significant in three networks (PA Tree, Erdos-Renyi, and Dense

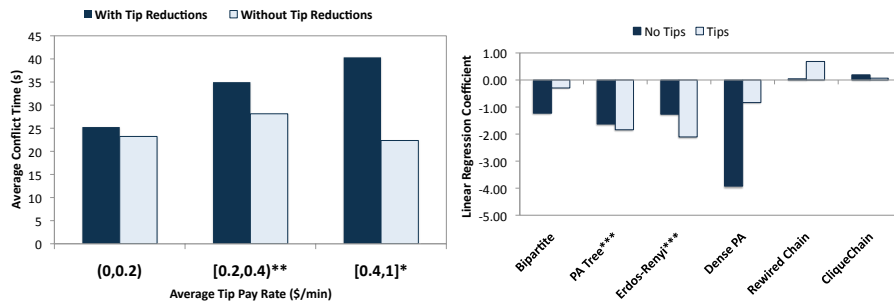


Fig. 6. Left: Average time in conflict at similar average tip pay rates with (black) and without (white) a negative tip change. Right: The connection (linear regression coefficient) between degree and income by network class for the no-tips and tips settings.

PA) in the no-tips setting and in only the first two in the tips setting. This is not too surprising: all the preferential attachment networks exhibit relatively low degree variation, and in rewired and clique chain graphs it is even less.

Nodes with a high degree spent considerably less time as King (correlation is -0.37 in the no-tips and -0.4 in the tips settings, both with $P < 0.001$). This finding is consistent across network classes. In contrast, the total time spent as King had significant *positive* correlation with wealth overall, 0.18 in the no-tips setting and 0.31 in the tips experiments ($P < 0.001$ in both). However, this conclusion is somewhat nuanced when dissected by network class: in 4 of the 7 network classes, the relationship between time spent as King and wealth is clearly positive in at least one of the game settings (no-tips or tips), but it is highly significant and *negative* in the two most highly clustered networks, rewired and clique chain. Thus, while generally being a King carries an advantage, it is more trouble than it's worth in highly clustered networks (presumably, because Kings face far too much conflict there from other neighbors vying for power).

While high degree nodes had a disadvantage, they were partially compensated for their handicap when tipping was allowed: the correlation between degree and tip income was 0.27 ($P < 0.001$); they naturally also dished out significantly less in tips to their neighbors (correlation between degree and tips paid was -0.23 with $P < 0.01$). Both these findings are consistent across network topologies.

4 Conclusion

One of our key observations is that allowing players to exchange tips substantially increases social welfare. Furthermore, we note that although conflict is clearly damaging to all parties, players systematically engage in it, although substantially less when tipping is allowed. We explain the impact of tipping on the amount of conflict between players by noting that tips equalize incomes between network neighbors. When players view their neighbors' income as unfairly higher than theirs, they engage in conflict, perhaps to punish the high earners. Greater equality in wealth therefore reduces the propensity to engage neighbors in conflict.

Since tip exchanges are pure transfers of wealth in our setting, classical economic theory would not anticipate any impact of tips on average profits. It is thus rather remarkable that tipping raises social welfare in our experiments. The positive welfare impact of tipping (and greater equality of wealth distribution) has considerable implications for policy, as it suggests that bridging income inequality may raise social welfare. Alternatively, our findings suggest that when compensation, resources, or tasks are distributed unequally, transfers of money or gifts may go a long way in alleviating interpersonal conflict.

References

1. George Akerlof and Robert J. Shiller. *Animal Spirits: How Human Psychology Drives the Economy and Why It Matters for Global Capitalism*. Princeton University Press, 2009.
2. George A. Akerlof and Janet L. Yellen. The fair wage-effort hypothesis and unemployment. *The Quarterly Journal of Economics*, 105(2):255–283, 1990.
3. Siegfried K. Berninghaus, Claudia Keser, and Bodo Vogt. Strategy choice and network effects. In A. Ockenfels and A. Sadrieh, editors, *The Selten School of Behavioral Economics*, chapter 8, pages 89–109. Springer-Verlag, 2010.
4. G.E. Bolton and A. Ockenfels. ERC: A theory of equity, reciprocity, and competition. *American Economic Review*, 90(1):166–193, 2000.
5. Roger Brown. *Social Psychology*. The Free Press, 2nd edition, 1986.
6. Colin F. Camerer. *Behavioral Game Theory*. Princeton University Press, 2003.
7. E. Fehr and K. Schmidt. Fairness, competition, and inequality. *Quarterly Journal of Economics*, 114(3):817–868, 1999.
8. Ernst Fehr and Simon Gächter. Cooperation and punishment in public goods experiments. *American Economic Review*, 90(4):980–994, 2000.
9. Drew Fudenberg and David K. Levine. *The Theory of Learning in Games*. The MIT Press, 1999.
10. M.R. Garey and D.S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman, 1979.
11. Stephen Judd and Michael Kearns. Behavioral experiments in networked trade. *Proceedings of the ACM Conference on Electronic Commerce*, 2008.
12. Stephen Judd, Michael Kearns, and Yevgeniy Vorobeychik. Behavioral dynamics and influence in networked coloring and consensus. *Proceedings of the National Academy of Sciences*, 107(34):14978–14982, 2010.
13. Michael Kearns, Stephen Judd, Jinsong Tan, and Jenn Wortman. Behavioral experiments on biased voting in networks. *Proceedings of the National Academy of Sciences*, 106(5):1347–1352, 2009.
14. Michael Kearns, Siddharth Suri, and Nick Montfort. An experimental study of the coloring problem on human subject networks. *Science*, 313(5788):824–827, 2006.
15. Tibor Neugebauer, Anders Poulsen, and Arthur Schram. Fairness and reciprocity in the HawkDove game. *Journal of Economic Behavior and Organization*, 66(2):243–250, 2008.
16. Mark Newman. *Networks: An Introduction*. Oxford University Press, 2010.
17. Ryan Oprea, Keith Henwood, and Daniel Friedman. Separating the Hawks from the Doves: Evidence from continuous time laboratory games. Working paper, 2010.
18. Matthew Rabin. Incorporating fairness into game theory and economics. *American Economic Review*, 83(5):1281–1302, 1993.