

A Behavioral Study of Bargaining in Social Networks

Tanmoy Chakraborty
University of Pennsylvania
tanmoy@seas.upenn.edu

Michael Kearns
University of Pennsylvania
mkearns@cis.upenn.edu

Stephen Judd
University of Pennsylvania
sjudd@seas.upenn.edu

Jinsong Tan
University of Pennsylvania
jinsong@seas.upenn.edu

ABSTRACT

We report on a series of highly controlled human subject experiments in networked bargaining. The basic interaction between two players is the decision of how to share a mutual payment; we extend this to situate the players in a network. Various theories predict, to different levels of uniqueness, what the shares will be. We analyze our experimental results from three points of view: social efficiency, nodal differences, and human differences; and contrast our behavioral results with the theories.

Categories and Subject Descriptors

J.4 [Social and Behavioral Sciences]: Economics

General Terms

Economics, Experimentation, Human Factors

Keywords

Network Economics, Behavioral Economics, Bargaining, Human Subject Experiments

1. INTRODUCTION

In recent years there has been much research on network-based models in game theory, in both the computer science and economics communities. Topics of attention include the effects of network topology on equilibrium properties, the price of anarchy in networking problems like selfish routing, game-theoretic models of network formation, equilibrium computation in networked settings, and many others. This large and growing literature has been almost exclusively theoretical, with few accompanying empirical or behavioral studies examining the relevance of the mathematical models to actual behavior.

In this paper, we report on a series of highly controlled human subject experiments in networked bargaining. Networked bargaining is modelled as follows: players are the

nodes of the networks, and each edge in the network represents some fixed amount of money that can be realized by its endpoints if they agree on how to split the amount. This agreement shall be referred to as *closing a deal*. In addition, there is a *deal limit* on each node, which is the maximum number of deals a player at that node is allowed to close, which could be less than its degree.

We were partly inspired by a long line of previous theoretical work which tried to relate wealth to network topology in bargaining settings [8, 9, 2, 14, 4, 13, 1, 5]. A notable feature of these theories is the prediction that there may be significant local variation in splits purely as a result of the imposed deal limits and structural asymmetries in the network. One can view our experiments as a test of human subjects' actual behavior at this game in a distributed setting using only local information. Our experiments are among the first and largest behavioral experiments on network effects in bargaining conducted to date.

We adopt many of the practices of behavioral game theory, which has tended to focus on two-player or small-population games rather than larger networked settings. In each of our experiments, three dozen human subjects simultaneously engage in one-to-one bargaining with partners defined by an exogenously imposed network. Our work continues a broader line of research in behavioral games on networks at the University of Pennsylvania [12, 10, 11]. Closest in spirit to the current work is that investigating networked exchange economies [10], but the experiments here and the theories underlying networked bargaining differ significantly from networked trading models.

In an extensive and diverse series of behavioral experiments, and the analysis of the resulting data, we address a wide range of fundamental questions, including: the relationships between degree, deal limits, and wealth; the effects of network topology on collective and individual performance; the effects of degree and deal limits on various notions of "bargaining power"; notions of "fairness" in deal splits; and many other topics.

The networks used are inspired from common models in social network theory, including preferential attachment graphs, and some specifically-tailored structures.

In all our experiments, the number of deals that were closed was above 85% of the maximum possible number. This is high enough to demonstrate real engagement, and low enough to demonstrate real tension in the designs.

Most of the deeper findings can be related to existing network bargaining theory. Although deals are often struck with unequal shares, more than one-third of the deals are

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

EC'10, June 7–11, 2010, Cambridge, Massachusetts, USA.
Copyright 2010 ACM 978-1-60558-822-3/10/06 ...\$10.00.

equally shared, thus indicating that people, while behaving as self-interested actors, also have an aversion towards inequality.

Network topologies have enough of an effect that they can be distinguished statistically via individual wealth levels and other measures. Higher degree, for example, tends to raise bargaining power while higher deal limits tend to decrease it. But while local topology affects bargains, invisible competition also affects it, even when the local topologies are indistinguishable. We find the expected effects of higher deal limits in the first neighborhood and higher degrees in the first and second neighborhoods, but neither degree distribution nor deal limit distribution is sufficient to determine the inequality of splits. In sum, there is a rich interaction between network and wealth that needs more study.

Other findings that speak to no existing theories but might provoke some new ones are the following:

- There is a positive correlation between inequality and social efficiency.
- Failures to agree on a split (as opposed to failures to find the best global trade configuration) form the greater part of missing efficiency.
- Social efficiency was higher when some uncertainty existed about a partner’s costs.

Finally, there are two curios that seem more about psychological dynamics than economics: People who are patient bargainers tend to make more money; and an incidental asymmetry in our protocol for closing a deal is correlated with a bias in the split.

In the ensuing sections, we review relevant networked bargaining theories, describe our experimental design and system, and present our results.

2. BACKGROUND

Networked bargaining with deal limits on the nodes, also known in the sociology literature as networked exchange with substitutable or negatively connected relations (eg. [4]), has been studied for decades. Several theoretical models have been designed to predict or propose how wealth should be divided [9, 14, 8, 16], and human subject experiments have been conducted on a few small graphs (up to 6 nodes) [6, 7, 16], albeit with different interfaces and mechanisms than ours. Some of the theoretical models are based on limited experimentation, along with simulated human behavior on slightly larger graphs [7]. A few models are based strongly on notions of game-theoretic rationality and are natural extensions of standard economic literature to social networks. Two models that belong to this class were introduced by Cook and Yamagishi [8] and by Braun and Gautschi [4]. We shall mainly focus on these two models.

The model given by Cook and Yamagishi, sometimes referred to as equidependence theory, is the most recognized theoretical model, and has received a lot of recent focus from the theoretical computer science community [13, 1]. Though Cook and Yamagishi [8] considered only unique exchange networks (that is, where each vertex may close only a single deal), the model is easily extendable to networks with varying deal limits. Every node is assumed to play strategically with selfish game-theoretic rationality. An *outcome* describes the division of wealth on various edges of the network. The *outside option* of a node is the highest offer it can rationally receive from any of its neighbors, such that closing that deal would benefit both parties, compared to the given

state. An outcome is said to be *stable* if every player’s earning is more than its outside option. Game-theoretic rationale suggests that an outcome should be stable if the players act in a myopically selfish manner. Cook and Yamagishi propose that the achieved outcome must be stable; moreover, they propose that the achieved outcome should be *balanced*, that is, two parties that close a deal should have equal additional benefit from this edge, where additional benefit is measured as the amount by which the earning of a player exceeds its outside option. Kleinberg and Tardos [13] showed that a stable and balanced outcome exists on all bipartite networks, but may not exist in all networks, and if it does, the closed deals in a stable outcome form a maximum matching. This equal division of surplus is stipulated by standard two-player bargaining solutions such as the Nash Bargaining Solution and Proportional Bargaining Solution, for players with linear utilities [15, 3].

Though a balanced outcome seems to be the most robust theoretical model, it has several drawbacks, first and foremost that it does not exist on even simple networks such as a triangle; and when it exists, there is a balanced outcome for every maximum matching in the network. This makes it computationally hard to even enumerate all the balanced outcomes in a network, and non-uniqueness reduces the predictive value of such a model. Another drawback is that the model often suggests that some edges will be shared so that one party gets an infinitesimal share, and the other party gets practically the entire amount. For example, a node that has at least two leaves (nodes of degree 1) as neighbors always ends up with maximum possible profit, due to competition between the leaves. However, even previous small-scale experiments [16] have suggested that such a phenomenon does not happen, and in our experiments, players rarely close a deal that extremely favors one of the players. Thus, when human subjects are involved, perfect local rationality seems to be an incorrect assumption.

The model given by Braun and Gautschi [4] defines a “bargaining power function” on nodes that depends only on the degree of the node and degrees of its neighbors. This function increases with increase in degree of the node, and decreases with increase in degrees of its neighbors, and is independent of all other network aspects. On each edge, the division of wealth, if a deal is made, is stipulated to be proportional to the bargaining power of the adjacent nodes. The bargaining power functions do not distinguish between different limits on nodes, and generally assume that relations are negatively connected: that is, for any given player, closing one deal reduces the maximum value that can be obtained from other edges deals. This makes the model quite inadequate as a predictor for our experiments. The other feature of this model is that network effects are quite local in nature, since even slightly distant properties such as the degrees of neighbors of neighbors do not have an effect on the bargaining power function. However, the model attempts to capture the notion that the earning of a player depends positively on its own degree and negatively on the degree of its neighbors. We test this notion on fairly large graphs for the first time, and we also show that the degrees of neighbors of neighbors do affect a node positively. Such alternating effects were predicted in previous theoretical models such as that by Markovsky et. al. [14], which said that odd length paths from a node enhance its earning, while even length paths reduce it.

The most significant set of previous experiments were done by Skvoretz and Willer [16], who conducted experiments on 6 small networks (each has at most 6 nodes), with only unit deal limits in 4 of them. They found that some common intuitions held true in those networks. For example, players who have deal limit one and multiple leaves as neighbors gets the bigger fraction of a closed deal, and that this fraction reduces if the limit of the player is raised. Among other results, we test such hypotheses extensively on much larger graphs with much more variance in their degree and limit distributions, and establish these hypotheses with very high statistical significance. Larger graphs also allow us to study the effects of network topology aspects that are more involved than the degree or limit of the player.

Recently, Chakraborty et. al. [5] designed an extension of the Cook-Yamagishi model, in the setting where there are no limits on the number of deals. In this case, the model predicts that all deals should be closed, and if players have linear utility (which is assumed in the Cook-Yamagishi model), all deals should be shared equally. Unequal splits may occur only if players have non-linear utility.

3. EXPERIMENTAL DESIGN

We designed 18 different experimental scenarios (consisting of specific choices of networks and arrangement of deal limits; each such scenario received 3 trials, for a total of 54 short experiments). These scenarios were based on 8 different graphs with a wide variety of details that is exemplified in Figure 1. The sole property they share is that they all have 36 nodes. We are thus casting our experimental nets wide here regarding network topology, as in much of our previous behavioral work. This section describes all the scenarios, at least at a high level.

The networks fall into 2 categories: regular graphs (to isolate and explore the effects of variations in deal limits), and irregular graphs (which contain an assortment of different degrees).

3.1 Irregular Graphs

We were interested in how bargaining behavior changes with changes in local network structure, and especially with differences in degree. Out of the huge space of such networks, we chose four. The first three we describe all had a common degree sequence, but differed in the way that nodes of each degree connected to nodes of other degrees. We generated a single degree sequence with a distribution that approximately follows a power law, and used it to build three graphs with different patterns of degree-to-degree profiles. We refer to these graphs as PL (for Power Law) graphs.

3.1.1 Power Law Graphs

Since we suspected that degree might have a large influence on bargaining power (to be confirmed below), it matters to the success of any node what the degrees are of other nodes they need to bargain with. Hence it was important to manipulate the *degrees of neighbors* as well.

By connecting nodes in different ways, we generated 3 graphs that differ in this manner but have the same basic degree distribution. In *PLP* (for Power Law Positive, indicating locally positive degree correlation) the high degree nodes are connected to other high-degree nodes. It models a world where nodes of different degrees are segregated from each other. In *PLN* (for Power Law Negative) the high de-

gree nodes are connected to the low-degree ones; it models a world where the connection-poor are likely to be ‘captivated’ by the connection-rich. And *PL0* has them all mixed together to disperse such phenomena. It models a world where nodes mingle freely with other types.

With each of the PL graphs above, we used each of the following 3 deal limit schemes to obtain $3 \times 3 = 9$ different scenarios. The first is the well-studied unique exchange situation (*uniq*): all nodes have deal limit 1. The other 2 are neither unique exchange nor unlimited, but represent two points in another large space of possibilities in between those notions. They are best thought of as having random deal limits drawn uniformly between 1 and the degree of the node. We call them *limA* and *limB*, and the difference is just that they are different randomizations.

3.1.2 Identical First Neighborhoods

The final irregular graph was designed specifically to test if structure outside the immediate neighborhood of a node would affect its behavior. The network used for this test has two sets of three identical nodes, which are colored blue and red in Figure 1d. Both sets have degree 6, and each of their neighbors have degree 7, so the local neighborhoods are indistinguishable in our GUI views. Any differences in behavior must be due to the second neighborhood or aspects even more distant. The second neighborhood of these two sets of nodes are drastically different; the second neighbors of the red nodes includes the 20 leaves while the second neighbors of the blues does not. This graph helped us identify the effects of second neighborhood when the first neighborhoods of two nodes were identical. We used it only with all nodes having deal limit 1. We refer to this scenario as *2ndHood*.

3.2 Regular Graphs

The 8 remaining scenarios were all based on regular graphs. This allowed us to test effects other than degree, like differing deal limits or large-scale market imbalances. One graph is the cycle shown in Figure 1. Four of them are identical tori with different deal limit schemes. Finally, three other graphs were used to observe the effects of a global supply imbalance, and are described in section 3.2.2.

3.2.1 Tori

The 4 tori are topologically the same as the 6×6 torus in Figure 1b, and are differentiated only through deal limits:

- Uniform Torus (*torUniq*): all nodes have deal limit 1.
- Checkerboard Torus (*torChkb*): all white nodes have deal limit 1, the others have deal limit 3.
- Torus Rows (*torRows*): alternating rows have deal limit 1 and deal limit 3.
- Torus Diamond (*torDiamnd*): Some vertices have deal limit 1 and some have deal limit 3. See Figure 8.

3.2.2 Imbalanced Supply Networks

The supply networks are 3 regular graphs which were designed to study the effect of a capacity issue which is not apparent at the node, but becomes apparent when contrasting the deal limits of two *groups* of nodes. Let the *external demand* of a group be the sum of deal limits of the nodes in the group minus the maximum number of deals that can be closed within the group.

In the supply networks, we defined the groups as the left group and right group as shown in Figure 2. All nodes have




Figure 1: Representatives of all the networks used. (a) the PA graph with positive homophily. (b) a torus; the edges running off the top and bottom denote wrap-around connections, as do those off the sides. (c) a simple cycle. (d) the 2ndHood graph for testing second neighborhood effects.

degree 4. All vertices in the left group have deal limit 2, and all vertices in the right group have deal limit 3. In each network, the right group has two different types of neighbors: those that belong to the right group, and those that belong to the left group. It is their differential treatment of the two types that was of interest. Nodes on the left have only one kind of neighbor; they exist just to set up the market conditions for those on the right.

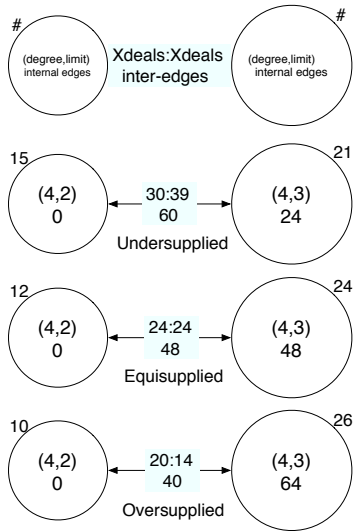


Figure 2: Metadesigns for graphs studying differential treatment of nodes under three global market conditions. The top line is a template for interpreting the others. Xdeals means external demand. The “supply” in the 3 names refers to the number deals that the right side wants versus the number available.

The three graphs share the fact that all deal limits are either 2 or 3, and the ones on the right have both types of neighbours. They are different in the ratios of external demands between the left and right groups; in the Undersupplied case the right nodes are somewhat starved for deals (seeking 39 when only 30 could be forthcoming), in Equisupplied they are just balanced, and in Oversupplied they have more offers than they can use.

4. SYSTEM OVERVIEW

Experiments were conducted using a distributed networked software system we have designed and built over the past several years for performing a series of behavioral network experiments on different games. This section briefly describes the user’s view of that system in our bargaining experiments.

Like most microeconomic exchange models, the model described in Section 2 does not specify an actual temporal mechanism by which bargaining occurs, but of course any behavioral study must choose and implement one. At each moment of our experimental system, and on each edge of the network, each human subject is able to express an *offer* that is visible to the subject’s neighbor on the other end of the edge. See Figure 3. The offer expresses the percentage of the benefit that a player is asking for. When the portions on either end of an edge add up to exactly 100%, one of the players is able to *close* the deal by pressing a special button. Individuals can always see the offers made to them by their neighbors, as well as some additional information (including the degrees and limits of their neighbors, and the current best offers available to their neighbors). When a deal is closed, or when one of the partners has used up his limit of deals, the relevant edge mechanisms are frozen and no further action is allowed on them. Every game is stopped after 60 seconds. Any money riding on deals not closed within that time is simply “left on the table”, i.e. the players never get it.

All communication takes place exclusively through this bargaining mechanism. Actions of a user are communicated to the central server, where information relevant to that action is recorded and communicated to the terminals of other users.

4.1 Human Subject Methodology

Our IRB-approved human subject methodology was similar to that of our previous experiments [12, 10, 11].

4.2 Session Overview

The main experimental session we shall study, which employed the network structure and deal limit scenarios described above, consisted of 3 trials each of the 18 scenarios described above, making 54 experiments in all. Each edge had a payment of \$2 available, and in the end approximately \$2500 was spent on subject payments. Unless mentioned




Figure 3: Screenshot of player’s interface for bargaining.

otherwise, this paper shall always be describing this session of experiments.

Prior to this main session, we ran a preliminary set of experiments that employed many but not all of the same network structures, but without any deal limits imposed. Some of these experiments also imposed “transaction costs” on vertices for closed deals. We will mention results from this earlier session and contrast them with those of our main session in a couple of places.

5. RESULTS

Our results come under three broad categories. The first is about collective performance and social efficiency. The second category examines questions about the differential fates of nodes, depending on their position in the networks and the deal limits they each had. The third category is about the general performance of humans summarizing behavior across all the games they played. This is an area that no economic theory attempts to cover.

5.1 Social Welfare

Humans were quite effective at playing these games, but they paid a surprising price for their refusal to close some deals.

To quantify how well humans did on this problem, we implemented a greedy algorithm for comparison. Given a graph and deal limits, it repeatedly draws (uniformly at random) an unclosed deal, both of whose endpoints have not already saturated their deal limit, and closes it, until there

are none left. To normalize both the human and greedy systems we divide by the *Maximum Social Welfare*, which is the maximum number of deals that can close in each network, subject to both topology and deal limits. The *social welfare* is the number actually closed, and the ratio between this and the max is the *social efficiency*.

The observed efficiencies are rendered by blue dots in Figure 4. In 6 of the networks (those below the diagonal), the humans did worse than the greedy algorithm. Full efficiency is rare in both systems. One might view this as the *behavioral price of anarchy* due to selfish players operating with only local information. The greedy algorithm obtained an average of 92.14% of the maximum welfare in our networks. In comparison, human subjects achieved an average social welfare of 92.10% of the maximum welfare when averaged over all 3 trials, a surprisingly similar figure.

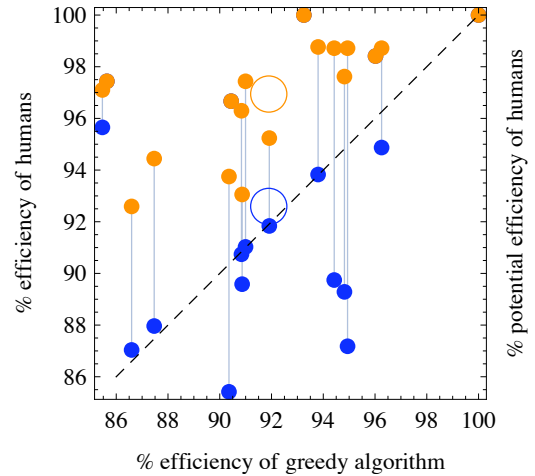


Figure 4: Scatterplot of greedy algorithm efficiency versus humans, one dot for each of the 18 scenarios (averaged over trials). Blue dots are what the humans actually achieved. Orange dots are the result of applying the greedy algorithm to the final state of human play, which is what the humans could have achieved without obstinacy. Vertical lines thus show the price of obstinacy. The dotted line indicates equality of the two scales. The open circles represent the average values over all scenarios.

There are two parts to this story, though, because solving these problems involves both selecting edges and closing deals on them. The greedy algorithm does not address the deal-closing issue and per force never leaves a potential deal unclosed; the humans often did. In 36 of the 54 experiments, the solution found by the human subjects was not even maximal – there were adjacent vertices that both could have closed another deal. Presumably this was because they simply could not agree on a split. However, the humans left the system in a state that could be improved post facto. We started the greedy algorithm in the final state the humans reached and allowed it to attempt to find more deals, thus producing a new state with no further unclosed deals. In all cases, this new state had a *higher* social efficiency than the greedy algorithm achieved alone. This is shown in the

orange dots of Figure 4. A line connects the human performance to the potential human performance, and we might dub this difference the *price of obstinacy*. In total, 7.9% of the money was “left on the table”, but 4.5% was due to obstinacy (more than half the lost value).

We conclude that the humans found better matchings in the graph, and hence their behavioral price of *anarchy* is lower (better) than the greedy algorithm. But due to their additional *obstinacy*, their overall performance was no better.

5.2 Nodal Differences

There was much evidence that nodal income depends on its deal limit, its degree, and properties of the non-local neighborhood.

5.2.1 Unequal Splits

Most theoretical models (for example, the Yamagishi-Cook model) that apply game-theoretic rationality to bargaining suggest that deals in some networks will be split in an unequal fashion. We will report the splits using their *inequality value* (Ineq), defined as the absolute difference between the two fractional shares. It ranges from 0 (equal sharing) to 1 (one player gets everything).

A total of 1271 deals were closed in all the 54 experiments, and 423 of them were split equally (inequality=0). But most were not split equally, every possible granular division was used for some splits, and 6 edges even had inequality=1 (which is surprising in itself since one partner gains nothing by signing the deal). The histogram in Figure 5 shows the inequality values. For comparison, we also show the histogram (in orange) from our preliminary session, which had no deal limits and produced an overwhelming portion of deals that split 50:50. The average inequality value over all games in our main session was 0.2097, which is a ratio of about 60:40. It thus seems clear that deal limits are invoking a significant increase in imbalanced splits.

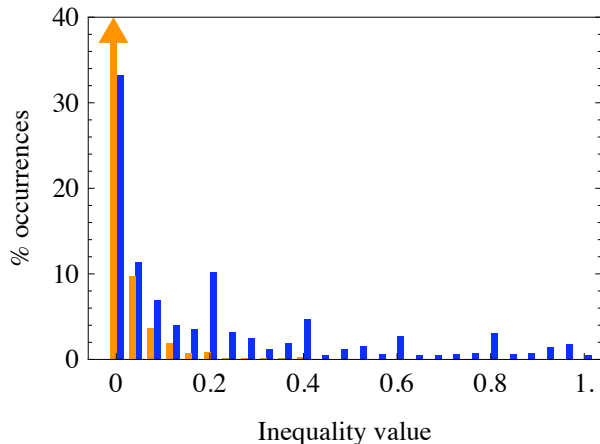


Figure 5: Histogram of deal splits in all games. The orange distribution was from our preliminary session, and the bar at 0 (equal shares) goes to 82%. The blue bars are from our main session, where we obtained a much greater spread of unequal splits.

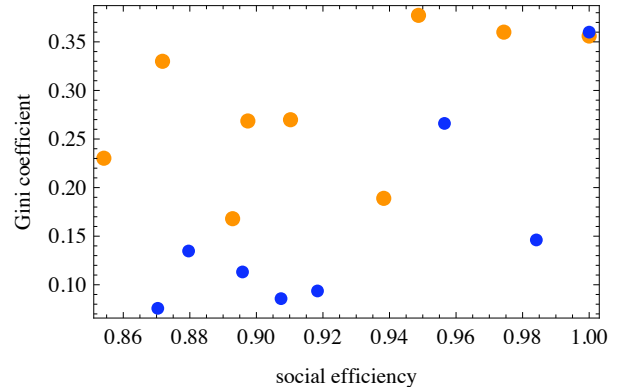


Figure 6: Average inequality value vs. Social Efficiency. The 9 PL networks are plotted in orange, and the 8 regular networks are plotted in blue.

5.2.2 Inequality and Efficiency

There is a significant correlation between average inequality value and the social efficiency achieved in each scenario — that is, *when the subjects collectively tolerate greater inequality of splits, social welfare improves*. These data are plotted in Fig. 6. The correlation coefficient is 0.52 with a confidence level of $p = .027$.

Interestingly, in an earlier session of experiments without deal limits, the same correlation is highly negative.

5.2.3 Degree Distribution and Equality

How well does degree distribution predict wealth distribution? We examined the PL networks to answer this. The average inequality value in closed deals is 0.23 over the 3 PLPuniq games (where nodes tend to be adjacent to nodes of similar degree), while it is 0.36 for both PLNuniq (where nodes tend to be adjacent to nodes of very different degree) and PL0uniq (where nodes tend to be adjacent to nodes of various degrees).

The inequality values of the PLPuniq experiments are less than the joint PL0uniq and PLNuniq outcomes with a one-sided $p < 0.02$. This indicates that *nodes that have an opportunity to bargain with at least one other of similar degree have more power than one that is forced to bargain only with higher-degree nodes*. (These are all unique-exchange games, so deal limit is not playing any distinguishing role.)

These three networks all have the same degree distribution. Hence, *degree distribution is not sufficient to predict inequality of wealth, even in unique-exchange networks*.

5.2.4 Deal Limit Distribution and Equality

A similar story holds for deal limit distribution. Even if the distribution of deal limits for two networks are identical, the experimental results can differ widely based on whether a node bargains with nodes of similar deal limits or differing deal limits.

In Torus-Uniform, all vertices have the same deal limit. In Torus-Rings, all vertices have two neighbors with the same deal limit (1 or 3) and two neighbors with a different deal

limit (3 or 1). In Torus-Checker, all nodes are bargaining with nodes of a different deal limit.

The average inequality values are 0.086, 0.13, and 0.36 for Torus-Uniform, Torus-Rows and Torus-Checker respectively. A means test shows these are all pairwise distinct with $p < .03$. These networks all have identical topologies. Thus, when network topology is not playing any distinguishing role, *if vertices bargain with vertices of similar deal limits, the deals are on more equal terms compared to when vertices with differing deal limits bargain.*

5.2.5 High Degree Confers Power

Over all closed deals in the PL graphs, the fractional take per closed deal of each node has a correlation of 0.47 with the degree of that node. If, to reduce the confound of differing deal limits, the study is confined to just the PL*uniq graphs, then the correlation coefficient is 0.59. Both these correlations are highly statistically significant. Thus, *bargaining power increases with the size of the local market*, at least in the setting where deal limits constrain behavior.

5.2.6 High Deal Limit Undermines Power

While higher degree confers bargaining power, higher deal limits had the opposite effect.

The Torus-Rows graph was designed specifically for testing the effect of deal limit on a node’s bargaining power. In this graph, all nodes are identical up to relabeling, but half of them have deal limit 1 and half have deal limit 3. So if there is any systematic difference between two nodes’ bargaining power, the difference in their limits can be the only explanation.

In the deals closed by limit-1 nodes, their mean fraction of the deal is 0.57. The limit-3 nodes obtained an average of 0.48. The difference was highly significant. (The two fractions do not add to unity because not all deals were between the two groups.) If only those deals between the two groups are considered, the fractions are 0.57 vs 0.43 and the difference is even more significant. The summary is that *a higher deal limit confers less bargaining power.*

5.2.7 Effect of Global External Demand

The supply networks were designed to study the effect of external demand. This property is not apparent at the node, but becomes apparent when contrasting the deal limits of two *groups* of nodes. In each supply network, the supply of deals from the left group (recall Figure 2) was manipulated to starve or overfeed the right group. How does the split of a deal depend on that relative supply?

In the Undersupplied case the right nodes must compete among themselves for the attention of nodes in the left side, so we might expect their shares to be smaller than the left side’s. In the Equisupplied case, the external demands are equal, so we might expect no differential in bargaining power. In the Oversupplied case, the left nodes must compete for deals from the right side, so we might expect their share of the deals to be smaller than the right.

Table 1 shows the results. There is a correlation of -0.19 ($p = 0.01$) between the external demand ratio and the deal share of the left nodes. The divisions favor the limit-2 nodes in all cases, consistent with the results of the previous section. However, that local property of relative limits is modulated by the global supply and demand ratio.

	external demand		avg shares	
	left	right	left	right
Undersupplied	30	39	0.57	0.43
Equisupplied	24	24	0.55	0.45
Oversupplied	20	14	0.52	0.48

Table 1: External demand imbalances shape bargaining results. The average splits shown are for edges between left and right nodes. Edges between right nodes have an average share of 0.5 by definition.

5.2.8 First Neighborhood Effects

We examined the three PL*uniq scenarios to find effects attributable to the degrees of first (one-hop) neighbors. For both nodes in all deals in the PL*uniq games, compute the fraction of the node’s take and the average of the degrees of its neighbors. The correlation between these quantities is -0.60 and is highly significant. Similar results occur when the data are restricted to just those nodes with some fixed degree.

The clear and consistent story in unique-exchange games is that *the share obtained decreases as the average degree of the neighboring nodes increases.*

The opposite story holds when the first neighbors have higher *deal limits*. We compared the 4-regular networks torRows and torChkb. In torRows, a vertex of deal limit 1 has two neighbors of deal limit 1 and two neighbors of deal limit 3, while in torChkb, a vertex of deal limit 1 has all four neighbors with deal limit 3. The mean share of the former was 0.57 while the latter obtained 0.68. The difference is statistically significant with $p = .0001$. *The bargaining power of a vertex is enhanced when neighboring vertices have higher deal limits.*

5.2.9 Second Neighborhood Effects

How does the network effect propagate beyond the immediate neighborhood? The 2ndHood structure has two sets of 3 nodes, each of which have identical degree and first neighborhood degrees. The results of the previous section will be mute about how these nodes fare.

However, the second (two-hop) neighborhood of these nodes are drastically different: the neighbors’ neighbors are leaves for 3 of them, and part of a clique for the other 3. The mean share of the first group was 0.347, the mean share of the second was 0.571, and the 2-sided p value was 0.027. *The bargaining power of a vertex is enhanced if its neighbors’ neighbors have higher degree.*

5.3 Comparison with Theoretical Models

We shall now point out some structural differences in solutions given by theoretical models and those found by human subjects. For our main session, where nodes have limits, we narrow our attention to the PL*uniq networks, since the Cook-Yamagishi model [8] was originally designed for unique exchange networks, and fails to make a stable prediction on the 2ndHood network. The model predicts that maximum social welfare (maximum matching) will be achieved on all the PL*uniq networks, which is rare in the experiments, as reported in Section 5.1.

Further, the model predicts that a node with at least two leaves (nodes of degree 1) as neighbors always ends up with

$1 - \epsilon$ fraction of a deal. This is due to myopic, rational competition between the leaves, where ϵ is the smallest non-zero amount that can be received by a node by signing a deal (this is the granularity of offers available in the GUI, and we let $\epsilon = 0.02$). Accordingly, the model predicts that there should be at least 30 such skewed deals in our experiments with the PL*uniq networks. In contrast, we find that there is 1 deal where one node gets 100%, 5 deals where one node gets 98%, and only 10 deals where one node gets more than 90%. Further, all but one of these deals are between a leaf node and a node of degree 5 or 6. This indicates that *extremely skewed deals are much rarer than what game-theoretic rationale suggests*, and is more likely when the degree differences are larger.

In our preliminary session, where nodes have no limits, unequal splits are rare, as reported in Figure 5. Chakraborty et. al. [5] designed a model for this setting. It predicts that all deals will be shared equally if the players ate the nodes have linear utility functions, and network effects may arise only due to non-linearity of player’s utility. So the results of the experiments can be explained in this model if we assume that *in our range of payoffs, the players have near-linear utility functions*. This is not very surprising, since a player can make only a few dollars in each experiment.

5.4 Human Subject Differences

Humans were randomly assigned to nodes in each experiment and randomly reassigned in each replication of a scenario. Hence none of the results above could be ascribed to human differences. However behavioral literature is replete with examples of how human subjects leave their stamp, and some traits emerge in our data too.

5.4.1 Patience

The correlation between the average time for each human to close a deal and the average gain from closed deals, aggregated over all deals in all games, is 0.6664 ($p = 0.00$). See data in Figure 7. Apparently, *patience pays off*.

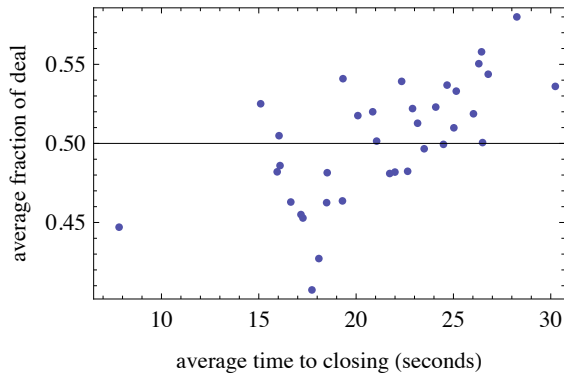


Figure 7: People with patience win bigger splits. There are 36 dots here, corresponding to the 36 human players.

5.4.2 Proposer vs. Acceptor

The user interface mechanism involved the following protocol for completing a deal: one player (*proposer*) makes

an offer, the other player (*acceptor*) accepts that offer by matching it, and then the proposer closes the deal. This was not designed to provoke any asymmetry, but was intended to avoid unintended closed deals due to accidental mouse-clicks. Nevertheless, by looking only at what the shares of the two parties were, can we say which party is more likely to have been the proposer? We find that indeed we can, and the party which gets the higher share is more likely to be the proposer. The mean proposer share across all experiments was 53.6%, and the acceptor share was 46.4%. The Kolmogorov-Smirnov test rejects the hypothesis that these shares come from identical distributions with $p < 10^{-14}$. Some psychological effect is clearly being expressed by this subtle asymmetry in protocol.

All possible split ratios in closed deals were at least once proposed by someone (and accepted by someone), with the sole exception of 0:100%. The six cases where all the money went to one player were all proposed by the high-share side. It may be “irrational” for someone to agree to get 0, but it would have been even odder to see someone *propose* that he get 0.

5.4.3 The Effect of Uncertainty in Costs

This last section is strictly about our preliminary experimental session, in which there were no deal limits imposed—so the limit on each node was effectively its degree. Here we found that social efficiency was higher when the players were simply *uncertain* about a particular detail regarding their neighbors.

In the latter half of the preliminary session, we imposed varying *transaction costs* on nodes, which a node must pay for every deal it closes. The first half of the experiments had no costs. Occasionally during the latter half, we quietly imposed zero cost on every vertex. This allows us to compare those games to the setting without costs.

This cost was specific to each vertex. Only the player at that vertex, but not its bargaining partners, knew how much this cost was. We varied the costs significantly, from 0 up to 40% of the value of the deal. This generated enough uncertainty that in the few instances where every vertex had zero cost, no one could infer the costs of his partner. This 0-cost setting can be directly compared to the basic non-costed setting where every player *knows* that there are no costs involved. Hence the two situations were distinguished only by a lack of certainty.

Players closed more deals in the (uncertain) 0-cost case than in the (known) no-cost case. The efficiency columns of table 5.4.3 show the fraction of possible deals that were closed in the two cases. The fraction went up in all 5 networks; the difference is significant with $p = .004$. Evidently, *the level of obstinacy rises when people know for certain that their partner has no costs*.

The average inequality values of the deals and the standard deviations are also shown in the table. We expected the splits to be more uneven in the zero-cost case, but no consistent story was found.

6. CONCLUSIONS

The background theory is not yet prepared to describe all the phenomena we have observed here. Some bargaining theory suggests one party to a deal might get an infinitesimally small share, but our mechanism does not allow this. Hence our results cannot be exactly matched, but the

	efficiency		average inequality		std. dev. of inequality	
	non	zero	non	zero	non	zero
PLP	0.85	0.96	0.012	0.01	0.033	0.037
PL0	0.84	0.93	0.009	0.009	0.025	0.029
PLN	0.72	0.93	0.027	0.012	0.057	0.051
cwc	0.84	0.95	0.015	0.023	0.035	0.076
2ndHood	0.84	0.97	0.014	0.009	0.040	0.044

Table 2: Social Efficiency and Inequality Values compared between the non-costed case and the zero-cost case. The cwc graph was a cycle with chords, shown in Appendix Figure 8. The others were as described for our main session.

scarcity of splits that are 98% or above seems to hint that the notion of “rationality” used by these theories needs to be adjusted. Other aspects of our results support theoretical models, notably the finding that phenomena at odd and even-length distances from a node alternately enhance and detract from the node’s earnings.

The findings peculiar to people –namely the prevalence of obstinacy, the value of patience, the effect of protocol in the closing of a deal, and the state of knowledge about the partners– are all in need of theoretical development. It seems these findings argue for the further need to integrate the fields of economics, game theory, sociology, psychology, and computer science.


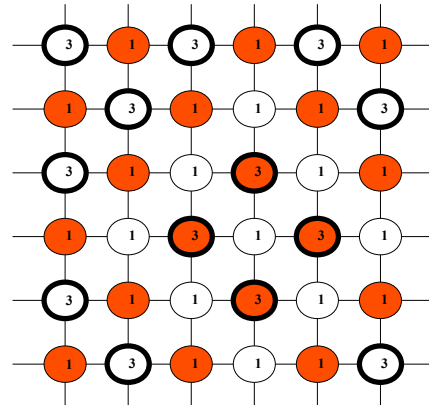
7. REFERENCES

- [1] Y. Azar, B. E. Birnbaum, L. E. Celis, N. R. Devanur, and Y. Peres. Convergence of local dynamics to balanced outcomes in exchange networks. In *FOCS*, pages 293–302, 2009.
- [2] E. J. Bienenstock and P. Bonacich. The core as a solution to exclusionary networks. *Social Networks*, 14:231–44, 1992.
- [3] K. Binmore. *Game Theory and the Social Contract, Volume 2: Just Playing*. The MIT Press, 1998.
- [4] N. Braun and T. Gautschi. A nash bargaining model for simple exchange networks. *Social Networks*, 28(1):1–23, 2006.
- [5] T. Chakraborty, M. Kearns, and S. Khanna. Network bargaining: algorithms and structural results. In *ACM Conference on Electronic Commerce*, pages 159–168, 2009.
- [6] K. S. Cook and R. M. Emerson. Power, equity, and commitment in exchange networks. *American Sociological Review*, 43 (5):721–739, 1978.
- [7] K. S. Cook, R. M. Emerson, M. R. Gillmore, and T. Yamagishi. The distribution of power in exchange networks: Theory and experimental results. *The American Journal of Sociology*, 89 (2):275–305, 1983.
- [8] K. S. Cook and T. Yamagishi. Power in exchange networks: A power-dependence formulation. *Social Networks*, 14:245–265, 1992.
- [9] N. Friedkin. An expected value model of social power: Predictions for selected exchange networks. *Social Networks*, 14:213–230, 1992.
- [10] S. Judd and M. Kearns. Behavioral experiments in networked trade. In *ACM Conference on Electronic Commerce (EC)*, pages 150–159, 2008.

- [11] M. Kearns, S. Judd, J. Tan, and J. Wortman. Behavioral experiments on biased voting in networks. *Proceedings of the National Academy of Sciences*, 106(5):1347–1352, 2009.
- [12] M. Kearns, S. Suri, and N. Montfort. An experimental study of the coloring problem on human subject networks. *Science*, 313(5788):824–827, 2006.
- [13] J. Kleinberg and E. Tardos. Balanced outcomes in social exchange networks. In *Symposium on Theory of Computing (STOC)*, pages 295–304, 2008.
- [14] B. Markovsky, J. Skvoretz, D. Willer, M. Lovaglia, and J. Erger. The seeds of weak power: An extension of network exchange theory. *American Sociological Review*, 58:197–209, 1993.
- [15] J. Nash. The bargaining problem. *Econometrica*, 18:155–162, 1950.
- [16] J. Skvoretz and D. Willer. Exclusion and power: a test of four theories of power in exchange networks. *American Sociological Review*, 58(6):801–818, 1993.

APPENDIX

a



b

Figure 8: (a) Diamond Torus, and (b) CWC (cycle with chords). CWC was used only in preliminary session.