Network Bargaining: Algorithms and Structural Results

Tanmoy Chakraborty
University of Pennsylvania
Philadelphia PA
tanmoy@cis.upenn.edu

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

Sanjeev Khanna
University of Pennsylvania
Philadelphia PA
sanjeev@cis.upenn.edu

ABSTRACT

We consider models for bargaining in social networks, in which players are represented by vertices and edges represent bilateral opportunities for deals between pairs of players. Each deal yields some fixed wealth if its two players can agree on how to divide it; otherwise it yields no wealth. In such a setting, Chakraborty and Kearns [5] introduced a simple axiomatic model that stipulates an equilibrium concept in which all players are rationally satisfied with their shares. We further explore that equilibrium concept here. In particular, we give an FPTAS to compute approximate equilibrium in bipartite graphs. We also show that equilibrium is not unique, and give conditions that ensure uniqueness on regular graphs. Finally, we explore the effect of network structure on solutions given by our model, using simulation methods and statistical analysis.

Categories and Subject Descriptors

F.2.2 [**Theory of Computation**]: Analysis of Algorithms and Problem Complexity; J.4 [**Social and Behavioral Sciences**]: Economics

General Terms

Algorithms, Economics, Experimentation, Theory

Keywords

Bargaining, Social Networks, Equilibrium, Approximation Algorithms

1. INTRODUCTION

Bargaining has been studied extensively in economics and sociology, both theoretically and experimentally. One setting that appears extensively in the literature is when there are only two parties negotiating a single deal. The deal yields a fixed total wealth if the two parties can agree on how to share or divide it; otherwise both parties receive nothing. Bargaining solution concepts provide predictions about how the wealth will be shared, or what division is "fair", which may depend on the player utility functions.

There are several bargaining solutions in economic theory, and here we shall focus on two of them: the *Nash Bargaining Solution (NBS)* and the *Proportional Bargaining Solution (PBS)*. Both of these solution concepts (and most others) predict that the division of wealth is a function of the *additional* utility (compared to some fixed "outside option" or alternative) each player receives by accepting the deal, which we shall refer to as *differential utility* (a formal definition is given shortly). NBS states that the division of wealth should maximize the product of the differential utilities of the two players, while PBS states that the division should maximize their minimum. When the players have increasing and continuous (in accrued wealth) utility functions, PBS simply states that the two players should have equal differential utility from the deal. We choose to focus on these two concepts because they are representatives of two broad classes of these solution concepts: NBS represents those solutions that do not allow direct comparison of utility across players, and are thus impervious to scaling of utility functions; whereas PBS represents those that permit such comparisons [3].

In this paper, we consider *multiparty, networked* generalizations of these classical bargaining frameworks and solution concepts. In our setting, players are represented by vertices and edges represent bilateral opportunities for deals between pairs of players. As in the two-player models, each deal yields some fixed wealth if its two players can agree on how to divide it; otherwise it yields no wealth. The yield of an edge is independent of that of other deals — however, network effects may arise due to the fact that the "outside option" of each player in considering one deal or edge is determined by the wealth they accrue from their *other* deals. Thus, for instance, a player with a concave utility for wealth and very high degree might have stronger bargaining power under certain solution concepts than a lower-degree neighbor.

A simple, intuitively justifiable, axiomatic model was introduced by Chakraborty and Kearns in [5] that states what division of wealth on every edge of the network rational players will consider acceptable. The proposed model is based on the assumption that the players are myopic, and act based on local information (about their network neighbors) only. Each player is rationally satisfied when she feels that she cannot get a greater share on any of her deals or edges if the deals between all other players remain unchanged. The

EC'09, July 6–10, 2009, Stanford, California, USA.

model stipulates that both endpoints of an edge are satisfied with the current division if and only if it satisfies a chosen two-player bargaining solution (e.g. NBS or PBS). The condition that all edges in the network are stable in this sense, when fixing the outcome of all other bilateral deals, gives rise to a network equilibrium concept for any given two-player bargaining solution.

Chakraborty and Kearns [5] showed that a PBS equilibrium exists on all networks if the players have increasing and continuous utility functions, and an NBS equilibrium exists on all networks if the utility functions are increasing, continuous and concave. Further, they showed some basic structural results about NBS equilibrium, and gave a polynomial time algorithm to compute approximate equilibrium on trees whose maximum degree is bounded by a constant.

1.1 Summary of Results

We first present a fully polynomial time approximation scheme (FPTAS) to compute approximate network bargaining equilibria in bipartite networks when the utility functions and the solution concept satisfy a natural condition that we call the *bargain monotonicity condition*. Bipartite networks are natural in many settings in which there are two distinct "types" of players — for instance, buyers and sellers of a good. The bargain monotonicity condition is satisfied by all concave utility functions in the PBS concept, and by natural utility function classes such as $x^p,\ 0 < p < 1$ and $\log(a + bx),\ a > 0, b > 0$ in the NBS concept. The algorithm can be viewed as iterating best-response dynamics on each edge under a particular schedule of updates; we show that for this schedule, the algorithm converges to (additive) ϵ-approximate equilibrium in time polynomial in ϵ^{-1} and the size of the input, if the values of wealth on the edges are polynomially bounded (for multiplicative approximation, we can handle arbitrary edge values). Whether the particular schedule for which we can prove fast convergence can be generalized is an interesting open problem.

We also perform simulation and statistical analyses of the effects of network structure on the wealth distribution at equilibrium for the two solution concepts, on networks randomly chosen from well-studied formation models (such as preferential attachment and Erdos-Renyi), and for a range of utility functions. Empirically we find that wealth of a vertex is highly correlated with degree, but degree alone doesn't determine wealth. We also find that bargaining power of a vertex, measured as the average share received by the vertex on all its edges, appears to increase with degree. Finally, we find stark differences between wealth distribution in PBS and NBS equilibrium. We find that the network effect is more pronounced in PBS equilibrium than in NBS equilibrium, which is manifested in two ways: first, the variation of bargaining power is larger in a PBS equilibrium than in an NBS equilibrium; and second, a higher number of edges have a highly skewed split in PBS than in NBS equilibrium. We also observe how network effects decrease as the utility functions approach linearity.

Finally, we show that neither PBS nor NBS equilibrium is unique on bargaining networks in general. Uniqueness is an important and preferred property, since it can serve as a prediction of how the wealth will be divided, as well as a measure of fair division. However, we unfortunately show non-uniqueness for the class of regular graphs with unit wealth on all edges, and the same utility function for all vertices. This class is interesting because every network in this class has one state which is both an NBS as well as a PBS equilibrium: the state where every edge is divided fifty-fifty. This state is also superficially fair, given the symmetry of the opportunities that the players have, as well as their behavior. Further, this state is the unique equilibrium when utility is linear, so it is natural to ask if the network has any "unfair" equilibrium for concave utility functions. On a positive note, we recognize conditions on utility $\mathcal{U}$ that makes PBS and NBS equilibrium in this class unique, and show that natural concave functions such as $\mathcal{U}(x) = x^p$ for any $0 < p < 1$ satisfies these conditions.

1.2 Related Work

The network setting that we consider falls under the heavily studied field of network exchange theory. In this area, many models consider the added restriction that every player has a limit on the number of deals she can get into, which is less than or equal to her degree in the network. Note that our setting corresponds exactly to this limited exchange setting when the limit for every player is equal to her degree. Several models have been proposed to predict what agreements the players will get into, and how will the wealth on these deals get shared (eg. [6, 16, 10, 2, 4]). Skvoretz and Willer [17] conducted human subject experiments to practically verify these theoretical predictions. Most of the focus has, in fact, been on unique exchanges, where every player can get into only one deal. Recently, Kleinberg and Tardos [15] analyzed the model given by Cook and Yamagishi [6] in the unique exchange setting, and found an elegant theoretical characterization, connecting the bargaining solutions to the theory of graph matchings.

All these models assume that all the players have the same behavior, and focus on the differences in bargaining power caused by network structure only. They also agree that these differences arise from the threat of exclusion, that is, a vertex can get into only a few deals and so has to ignore the offers of some of her neighbors; so she can ask for a better offer from her neighbor by threatening to get into a deal with some other neighbor instead. In the scenario where the limit on the number of deals is equal to degree for every vertex, there is no threat of exclusion, and all the models predict that there will be no network effect, and the wealth is divided into two equal shares on every edge.

However, it is important to note here that all these models implicitly or explicitly assume linear utility functions. Our model takes into account that players may have non-linear utility. In particular, we focus on increasing concave utility functions, that is, those with *diminishing marginal utility*. Diminishing marginal utility is a well-known phenomenon, and is also often used in financial theory to capture risk aversion. Our model agrees with these previous models when the utility functions are linear. In fact, when the utility functions are linear, the concept of PBS equilibrium is identical to the equi-dependence theory of Cook and Yamagishi [6]. However, interesting network effects appear when the players have concave utility functions, even if all players have the same utility function.

Another group of concepts related to our setting arise in coalitional games, such as Shapley value, core, kernel, nucleolus and bargaining sets. These concepts often involve the ability of players to form arbitrary coalitions, which implicitly assumes that the players have information about the entire network, and are not acting myopically based on local information. Thus, these concepts assume that players can solve problems that are computationally difficult. In fact, Deng and Papadimitriou [8] showed that many of these concepts are computationally hard, and suggested that a solution concept is appropriate only if it is efficiently computable. In strong contrast, our model only expects simple selfish behavior from the players. Perhaps not unrelated to this aspect of our model, we shall show that equilibria in our concept is computable in polynomial time on bipartite graphs, and natural heuristics perform well in our simulations on random graphs.

The rest of the paper is organized as follows: In Section 2, we define our model and all the basic concepts, along with some basic lemmas. In Section 3, we present our algorithm and its analysis. In Section 4, we report our simulations and statistical analysis. In Section 5, we present our results on uniqueness of equilibria. In Section 6, we point out that our model can be extended to general forms of utility functions, and our algorithm applies in this extended model as well. We conclude with some open problems in Section 7.

2. PRELIMINARIES

A *bargaining network* is an undirected graph $G(V, E)$ with a set V of n vertices and a set E of m edges, where vertices represent player and edges represent possible bilateral trade deals. There is a positive value $c(e)$ associated with each edge $e \in E$, which is the wealth on that edge. There is also a utility function $\mathcal{U}_v$ for each player v. The utility functions are all represented succinctly and are computable in polynomial time.

Let $e_1, e_2 \dots e_m$ be an arbitrary ordering of the edges in E, where e_i has endpoints u_i and v_i, $\forall i = 1, 2 \dots m$. A *state of the bargaining network* is described by the division of wealth on each edge of the graph. Let $x(u_i, e_i)$ and $x(v_i, e_i)$ denote the wealth u_i and v_i receive from the agreement on the edge e_i, respectively. Note that $x(v_i, e_i) = c(e_i) - x(u_i, e_i)$. We shall represent a state of the bargaining network as a vector $s = (s_1, s_2 \dots s_m) \in \mathbb{R}^m$ such that $s_i = x(u_i, e_i)$. Note that s uniquely determines the division of wealth on all edges.

Let $s \in \mathbb{R}^m$ be a state of the bargaining network G. For any vertex u and any edge e incident on u, let $\gamma_s(u)$ denote the total wealth of a vertex u from all its deals with its neighbors. Let $x_s(u, e)$ denote the wealth u gets from the agreement on edge e. Let $\alpha_s(u, e) = \gamma_s(u) - x_s(u, e)$ be the wealth u receives from all its deals except that on e. We say that $\alpha_s(u, e)$ and $\alpha_s(v, e)$ are the *outside options* with respect to the edge e for u and v respectively, that is, the amount each of them receives if no agreement is reached on the deal on e.

Definition 2.1 (Differential Utility) *Let s be any state of the bargaining network. Let x be the wealth of u from the deal on $e = (u, v)$. Then, the* differential utility *of u from this deal is $\Delta_u(x) = \mathcal{U}_u(\alpha_s(u, e) + x) - \mathcal{U}_u(\alpha_s(u, e))$, and the* differential utility *of v from this deal is $\Delta_v(c(e) - x) = \mathcal{U}_v(\alpha_s(v, e) + c(e) - x) - \mathcal{U}_v(\alpha_s(v, e))$.*

We shall drop the subscript s if the state is clear from the context.

Definition 2.2 *Let s be any state of the bargaining network. Define $y_s(u, e)$ to be the wealth u would get on the edge $e = (u, v)$ if it is renegotiated (according to some two-party solution), the wealth divisions on all other edges remaining unchanged. Also define* $\texttt{change}(s, e) = |x_s(u, e) - y_s(u, e)|$.

2.1 Proportional Bargaining Solution (PBS)

We say that the allocation on an edge $e = (u, v)$ with value c satisfies the *Proportional Bargaining Solution* (PBS) condition if it maximizes the function $W_P(x) = \min\{\Delta_u(x), \Delta_v(c(e) - x)\}$ where x denotes the allocation to u. Thus if an edge e is renegotiated according to the *Proportional Bargaining Solution* (PBS), then $y_s(u, e) = \arg\max_{0 \le x \le c} W_P(x)$ and $y_s(v, e) = c(e) - y_s(u, e)$. Note that $y_s(u, e)$ is simply a function of the two values $\alpha_s(u, e)$ and $\alpha_s(v, e)$, along with the utility functions of u and v.

The following lemma gives a simpler equivalent condition for PBS when the utility functions are increasing and continuous, and is applicable to the two-party setting as well.

Lemma 2.1 *If the utility functions of all vertices are increasing and continuous, then for any edge $e = (u, v)$, the PBS condition reduces to the condition $\Delta_u(x) = \Delta_v(c(e) - x)$, that is, the condition of equal differential utility, and there is a unique solution x satisfying this condition.*

2.2 Nash Bargaining Solution (NBS)

We say that the allocation on an edge $e = (u, v)$ satisfies the *Proportional Bargaining Solution* (PBS) condition if it maximizes the function $W_N(x) = \Delta_u(x)\Delta_v(c(e) - x)$ where x denotes the allocation to u. Thus if an edge e is renegotiated according to PBS, then $y_s(u, e) = \arg\max_{0 \le x \le c} W_N(x)$ and $y_s(v, e) = c(e) - y_s(u, e)$. Note that $y_s(u, e)$ is simply a function of the two values $\alpha_s(u, e)$ and $\alpha_s(v, e)$, along with the utility functions of u and v.

If e is renegotiated according to the *Nash Bargaining Solution* (NBS), then $y_s(u, e)$ is a value $0 \le x \le c$ such that the NBS condition is satisfied, that is, the function $W_N(x) = \Delta_u(x)\Delta_v(c(e) - x)$ is maximized.

The following lemma gives a simpler equivalent condition for NBS when the utility functions are increasing, concave and twice differentiable, and is applicable to the two-party setting as well.

Lemma 2.2 *If the utility functions of all vertices are increasing, concave and twice differentiable, then for any edge $e = (u, v)$, the NBS condition reduces to the condition $\frac{\Delta_u(x)}{\Delta'_u(x)} = -\frac{\Delta_v(c(e)-x)}{\Delta'_v(c(e)-x)}$, that is, the condition of equal differential utility, and there is a unique solution x satisfying this condition. Moreover, let $Q_u(x) = \frac{\Delta_u(x)}{\Delta'_u(x)}$, and let $R_v(x) = -\frac{\Delta_v(c(e)-x)}{\Delta'_v(c(e)-x)}$. Then $Q_u(x)$ is increasing, $R_v(x)$ is decreasing, and $Q_u(x) - R_v(x)$ has a unique zero in $[0, c(e)]$.*

2.3 Stability and Equilibrium

Definition 2.3 (Exact Stability and Equilibirum) *We say that an edge e is* stable *in a state s if renegotiating e does not change the division of wealth on e, that is,* $\texttt{change}(s, e) = 0$. *We say that a state s is an equilibrium if all edges are stable.*

We also study two notions of approximation, namely, additive and multiplicative approximations, as defined below.

Definition 2.4 (Additive ϵ-Stability and Equilibrium) *We say that an edge e is* ϵ-stable *in the additive sense in a state s if* $\texttt{change}(s, e) < \epsilon$. *We say that s is an* additive ϵ-approximate equilibrium *if all edges are additive ϵ-stable.*

Definition 2.5 (Multiplicative ϵ-Stability and Equilibrium) *We say that an edge e is* ϵ-stable *in the multiplicative sense in a state s if $|y_s(u,e) - x_s(u,e)| < \epsilon x_s(u,e)$ and $|y_s(v,e) - x_s(v,e)| < \epsilon x_s(v,e)$. We say that s is a* multiplicative ϵ-approximate equilibrium *if all edges are multiplicative ϵ-stable.*

In this paper, an approximate equilibrium will refer to additive approximation, unless specified otherwise.

We refer to an equilibrium as an *NBS equilibrium* if the renegotiations satisfy the NBS condition. We refer to the equilibrium as a *PBS equilibrium* if the renegotiations satisfy the PBS condition.

2.4 Bargaining Concepts as Nash Equilibria

The bargaining solutions may also be viewed as pure Nash equilibria of certain games. Each edge is a player, and an edge e has a strategy space $[0, c(e)]$. Strategy of an edge corresponds to the division of wealth on it. Let $e = (u, v)$ be an edge playing strategy x. If the payoff of e is $\Delta_u(x)\Delta_v(c(e) - x)$, and each edge wishes to maximize its own payoff, then the pure Nash equilibria of this game are exactly the NBS equilibria of the network. Similarly, if the payoff is instead defined to be $\min\{\Delta_u(x), \Delta_v(c(e) - x)\}$, then the pure Nash equilibria of the game coincides with PBS equilibria.

Thus, updating an edge corresponds to an edge playing a best response move in the corresponding game. As with all pure Nash equilibria concepts, we thus have a natural heuristic that gives an equilibrium if it terminates: start from an arbitrary state, and then update unstable edges repeatedly till all edges are stable. This heuristic is called *best response dynamics*. It is worth noting that approximate equilibria of this game does not necessarily coincide with approximate equilibria of the bargaining network with the same approximation factor.

3. COMPUTING EQUILIBRIA

We will now design a fully-polynomial time approximate scheme for computing approximate additive and multiplicative equilibria in bipartite networks, provided the bargain monotonicity condition below is satisfied.

3.1 The Bargain Monotonicity Condition

Condition 3.1 (Bargain Monotonicity Condition) *An instance of the bargaining problem satisfies the* bargain monotonicity condition *with respect to a solution concept (PBS, NBS) if for any edge $e = (u, v)$ and a pair of states s and s' of the bargaining network such that e is stable in both s and s' with respect to the solution concept, whenever $\alpha_{s'}(u, e) \geq \alpha_s(u, e)$ and $\alpha_{s'}(v, e) \leq \alpha_s(v, e)$, then $x_{s'}(u, e) \geq x_s(u, e)$.*

The above condition states that on any edge (u, v), if the outside options of u increases while that of v decreases, then u claims a higher share of wealth on this edge when it is renegotiated. Note that Condition 3.1 is essentially a two-player condition that has no dependence on the network itself and merely depends on the outside options of the players.

The bargain monotonicity condition seems like a natural condition that a negotiation between two selfish players can be expected to satisfy. It is a condition that depends on the bargaining network as well as the solution concept. The condition is indeed satisfied by the PBS solution concept on all networks whenever the utility functions are concave and increasing. For NBS, however, concavity and monotonicity of utility functions is not sufficient for satisfying this condition. We will instead identify a stronger property of utility functions that is necessary and sufficient for satisfying the bargain monotonicity condition on all networks. This stronger property is satisfied by several natural classes of utility functions including x^p, $0 < p < 1$ and $\log(a + bx)$ where $a, b > 0$.

Lemma 3.1 *PBS solutions satisfy Condition 3.1 on all networks where the utility functions of all the players are concave and increasing.*

PROOF. Let $p = x_s(u, e)$ and $q = x_s(v, e)$. Consider the state s'' derived from s with the sole modification that u gets p and v gets q on e. Then, the differential utility of u from e in s is at most that in s'', since the function $\mathcal{U}(z+p) - \mathcal{U}(z)$ is decreasing in z when $\mathcal{U}$ is concave (this is precisely equivalent to *diminishing marginal utility*). By a symmetric argument, the differential utility of v from e in s is at least that in s'. Thus, the differential utility of e to u is at most that of e to v in s''. So by Lemma 2.1, u must get at least p on the edge e in s' to ensure that e is stable. $\square$

We now focus on identifying utility functions where the NBS concept will satisfy Condition 3.1 on all networks. Let $\mathcal{U}$ be the utility function of any vertex u. In light of Lemma 2.2, it is clear that $R(\alpha, x) = \frac{\mathcal{U}(\alpha+x)-\mathcal{U}(\alpha)}{\mathcal{U}'(\alpha+x)}$ must be a non-increasing function of α. If not, then there exists some positive α_1 and $\alpha_2 > \alpha_1$, such that $R(\alpha_1, x) < R(\alpha_2, x)$. It is easy to create a network with two states s and s' where u have outside options α_1 and α_2 respectively, while one of its neighbor v have the same outside option in both states. Then, the balanced outcome on the edge (u, v) gives a greater share to u in s than in s', thus contradicting Condition 3.1. The utility functions of all players must satisfy this property, and it is sufficient as well. The above discussion is captured in the following lemma.

Lemma 3.2 *Let χ be the family of all concave, increasing, and twice differentiable utility functions $\mathcal{U} \in \chi$, such that $R(\alpha, x) = \frac{\mathcal{U}(\alpha+x)-\mathcal{U}(\alpha)}{\mathcal{U}'(\alpha+x)}$ is a non-increasing function of α for all $\alpha > 0$ and $x > 0$. Then every network, where all players have utility functions from χ, satisfies the bargain monotonicity condition for the NBS concept.*

Now, suppose that $\mathcal{U}$ is concave, increasing and twice differentiable at all positive values. Simplifying the equation, we have

$$\frac{\mathrm{d}}{\mathrm{d}\alpha} \frac{\mathcal{U}(\alpha + x) - \mathcal{U}(\alpha)}{\mathcal{U}'(\alpha + x)} \leq 0$$

$$\Leftrightarrow \frac{\mathcal{U}(\alpha + x) - \mathcal{U}(\alpha)}{\mathcal{U}'(\alpha + x)} \leq \frac{\mathcal{U}'(\alpha + x) - \mathcal{U}'(\alpha)}{\mathcal{U}''(\alpha + x)}$$

(since $\mathcal{U}'(\alpha + x) > 0$, and $\mathcal{U}''(\alpha + x) \leq 0$)

It can be easily verified that natural utility functions such as $x^p,\ 0 < p < 1$ and $f + \log(a + bx),\ a > 0, b > 0, f \geq 0$ belong to χ, and so Lemma 3.2 applies.

Not all concave utility functions satisfy the equation above. In particular, we construct a concave, increasing and twice differentiable function $\mathcal{U}$ that violates the equation. The key is to construct a sharp change in marginal utility. We achieve this by making $|\mathcal{U}''(\alpha + x)|$ very large compared to the other expressions in the equation. We describe the utility function by defining $\mathcal{U}'$. Let $\mathcal{U}'(x) = 1$ for $0 < x < 1$, and $\mathcal{U}'(x) = 1/2$ for $x > 1.01$. For $1 \leq x \leq 1.01$, $\mathcal{U}'(x)$ decreases from 1 to 1/2 smoothly, so that $\mathcal{U}'$ is differentiable everywhere, and we can also ensure that $\mathcal{U}''(1.005) < -50$, which is the average slope of $\mathcal{U}'$ in the range $[1, 1.01]$.

Now, in the above equation, let $\alpha = 0.5$ and $x = 0.505$. Note that the left-hand-side of the equation is more than $\frac{x\mathcal{U}'(\alpha+x)}{\mathcal{U}'(\alpha+x)} = x > 0.5$, while the right hand side is less than $1/|\mathcal{U}''(\alpha + x)| < 1/50$, thus violating the equation. So we conclude that if the marginal utility of a player changes abruptly, Condition 3.1 may be violated. The above discussion is summarized in the following lemma.

Lemma 3.3 *There exists a bargaining network where all the players have concave utility functions, such that Condition 3.1 is not satisfied in the NBS concept.*

3.2 Algorithmic Results

To be used as a subroutine in our algorithm, we define an `Update` oracle, that takes an edge $e = (u, v)$ of the network as input. The oracle is called when our algorithm shall have a *current* state s of the bargaining network, and the oracle shall renegotiate the edge e according to the 2-player bargaining solution we use, and modify the division of wealth on e only, to change the state to s', so that e becomes stable. The computation of the `Update` oracle depends solely on the outside options of u and v with respect to e, that is, $\alpha_s(u, e)$ and $\alpha_s(v, e)$, and the utility functions of the u and v. We assume that the oracle also knows the input to the problem, that is, the social network and the utility functions of the players, as well as whether the goal is to compute an NBS or a PBS equilibrium.

The `Update` oracle essentially performs an improvement step of the best response dynamics in the game played by edges that was described in Section 2.4. Since the bargaining equilibrium concepts are pure Nash equilibria of this game, it is natural to wonder if a sequence of updates starting from a random state of the bargaining network converges to equilibrium. We do not know if all sequences converge, though our simulations suggest that random sequences converge to approximate equilibrium on random networks. As we shall show, there exists a sequence that converges to additive ϵ-approximate equilibrium in number of steps that is polynomial in ϵ^{-1} and the number of edges, in bipartite graphs. The property of graphs that is critically used is that it has no odd cycles, which is an equivalent characterization of bipartite graphs.

Theorem 1 *If the bargaining network and the solution concept satisfies the bargain monotonicity condition, then from any start state s, there is a polynomial-time computable sequence of at most $O(m^2 c_{\max}/\epsilon)$ edge updates that converge to an additive ϵ-approximate bargaining equilibrium, where m is the number of edges. Thus we have an FPTAS if the* `Update` *oracle runs in polynomial time.*

> **Input**: Edge e
>
> Modify current state s to a new state s' such that $x_{s'}(u, e) = y_s(u, e)$ and $x_{s'}(v, e) = y_s(v, e)$, and s' matches s on all edges except e;

Function: `Oracle Update(`e`)`

To get an FPTAS for computing a multiplicative approximate equilibrium (which is an algorithm almost identical to that for the additive approximation), we shall need the `Update` oracle to further satisfy a basic condition, which is essentially that whenever an edge is updated, none of its endpoints get too small a share.

Condition 3.2 (Polynomially Bounded Updates Condition) *There exists a constant $r > 0$ such that if the outside options of the endpoints of an edge $e = (u, v)$ are at most α, then $y_s(u, e) \geq \frac{1}{\alpha^r}$ and $y_s(v, e) \geq \frac{1}{\alpha^r}$ for all $\alpha \geq 0$.*

Theorem 2 *If the bargaining network (and the solution concept) satisfies the bargain monotonicity condition and the polynomially bounded updates condition, then from any start state s, there is a polynomial-time computable sequence of at most $O(m^2 \log nc_{\max}/\epsilon)$ edge updates that converge to a multiplicative $(1+\epsilon)$-approximate equilibrium, where $c_{\max}$ is the maximum value of any edge in the network, n is the number of vertices and m is the number of edges. Thus we have an FPTAS if the* `Update` *oracle runs in polynomial time.*

Before describing the algorithm, we shortly dwell on when and how is the response of the `Update` oracle computable in polynomial time. It is true for all increasing and continuous utility function in the PBS concept, and also true for all increasing concave functions in the NBS setting. This is because by an application of Lemma 2.1 and Lemma 2.2 respectively, the problem reduces to being given two functions of the same variable, one increasing and the other decreasing, and being asked to find a value of the variable where the two functions are equal. This can be done with exponential accuracy in polynomial time using a binary search process. It is easy to absorb this exponentially small error in the update oracle into the approximation factor of the equilibrium, so we will neglect it.

We now proceed to describe our algorithm.

3.3 The Algorithm

Let s be the current bargaining state, describing division of wealth on each edge. Algorithm 1 describes our algorithm to compute an ϵ-approximate equilibrium. For either additive or multiplicative approximation, note that the corresponding definition for approximate stability of an edge should be used in the algorithm to decide if an edge is ϵ-stable.

It is fairly easy to see that when Algorithm 1 terminates, the final state of the bargaining network is an ϵ-approximate equilibrium. The outer `while` loop implies that algorithm can terminate only when all edges are colored black. Moreover, the inner `while` loop can terminate only when all black edges become ϵ-stable. Thus in the last repetition of the

Input: Bargaining network G and an oracle access to the `Update` function
Output: An ϵ-approximate equilibrium
Initialize to an arbitrary state s;
Color all edges WHITE;
while *there exists a* WHITE *edge* e **do**
  Color e BLACK;
  while *there exists a* BLACK *edge* e' *that is not* ϵ*-stable in the current state* **do**
    `Update` (e');
Output current state;

Algorithm 1: An FPTAS for computing an ϵ-approximate equilibrium

outer loop, the last white edge gets colored black, and then the inner `while` loop ensures that the algorithm terminates only when all the black edges (which is the entire network now) are ϵ-stable, and thus the current state is an ϵ-equilibrium. So we only need to argue the termination and running time of the algorithm.

Since the inner loop of the algorithm terminates only when all the black edges are ϵ-stable, so at the beginning of the next inner loop, only the new black edge e may not be ϵ-stable. If it is not, the loop terminates without a single call to the oracle. However, if e is unstable, then it is the first edge to be relaxed in this execution of the loop, and the influence of this update now travels along the black edges.

Let $e = (u, v)$, and suppose updating e caused u to receive more wealth from the deal on e, than it was receiving just before the update. In such a case, we say that the update *favors* u. Noting that G is a bipartite graph, we label every vertex in the same partition as u as $+$, and the vertices in the other partition as $-$. This labeling is only for the sake of analysis, and may be different for distinct executions of the inner loop. Note that every edge is between a vertex labeled $+$ and a vertex labeled $-$. Lemma 3.4 is crucial, and Theorem 1 follows from it. The proof of 3.4 depends on Condition 3.1 being satisfied.

Lemma 3.4 *If Bargaining Monotonicity Condition holds, then during the execution of the inner loop, whenever an edge $e' = (u', v')$ is updated, such that u' is labeled $+$, the update favors u'.*

PROOF. We shall prove this statement by induction. By definition, the statement holds for the first step, that of updating e. For the inductive step, suppose that the statement has been true for the first $i-1$ update steps, $i \geq 2$. Consider the i^{th} update step, where the edge updated is $e' = (u', v')$. Since the beginning of the loop, whenever an edge incident on u' has been updated, by the induction hypothesis, u' was favored as it is labeled $+$. Note that since e' is a black edge, it was ϵ-stable at the beginning of this loop. So, since the last time that e' was ϵ-stable, all the updates have only increased the outside option of u' with respect to e'. By a similar argument, the outside option has gone down or stayed the same for v'. Thus Condition 3.1 implies that u' is favored in this update step, and the inductive proof is complete. □

Proof of Theorem 1. If we only update edges that are not ϵ-stable in the additive sense, then since every update increases the wealth of the favored vertex by at least ϵ, and since the wealth from an edge e' to any of its endpoint cannot exceed $c_{\max}$, the edge shall not be updated more than $c_{\max}/\epsilon$ times in one iteration of the inner loop. Finally, in each repetition of the inner `while` loop, at most m distinct edges are updated, and the loop itself is repeated m times. This completes the proof of Theorem 1. □

Proof of Theorem 2. Now, suppose we only update edges that are not ϵ-stable in the multiplicative sense. Assume that the Polynomially Bounded Updates Condition is true. Then, whenever an edge has been updated at least once before, both of its endpoints receive a share whose inverse is polynomially bounded in $nc_{\max}$, which is an upper bound on all outside options. Either of these shares can at most go up to $c_{\max}$ within a single iteration of the outer loop, and on every update it goes up (or down) by a factor of at least $(1+\epsilon)$ (or $(1-\epsilon)$) so the edge cannot be updated more than $O(\epsilon^{-1}\log(nc_{\max}))$ times. Again, in each repetition of the inner `while` loop, at most m distinct edges are updated, and the loop itself is repeated m times. This completes the proof of Theorem 2. □

4. SIMULATION STUDIES

In this section we undertake simulation studies of the effects of network structure on bargaining equilibria. Previous theoretical work [5] has shed partial light on how structure influences wealth in NBS equilibrium. In particular, for concave utilities such as $\sqrt{x}$, an edge between two vertices of high degree (both degrees exceeding d) is shared almost equally, with both parties getting a fraction $\frac{1}{2} \pm \frac{1}{d}$ of the deal. While a number of works have strived to quantify relationships between network structure and various game-theoretic equilibria [15, 11, 7, 4], precise characterizations are rare. We thus turn to an alternative approach, which is that of empirically investigating the structure-equilibrium relationship in networks randomly generated from well-studied stochastic formation models [13, 12, 14, 17]. we are particularly interested in how a player's position in the network influences her "bargaining power" [15, 4, 13, 17].

4.1 Methodology

The broad methodology we followed was to (a) generate many random networks from specific stochastic formation models (namely, preferential attachment [1] and Erdos-Renyi [9]), (b) compute bargaining equilibria on these networks by running best-response dynamics until convergence, and (c) perform statistical analyses relating structural properties of the network to equilibrium properties.

For the best response dynamics, we start from a random state, and then repeatedly pick any edge that is not ϵ-stable and update it, until all edges are ϵ-stable, for $\epsilon = 0.001$. In all our simulations, all edges in the network have unit wealth. Also, in all our simulations, we imposed the same utility function $\mathcal{U}$ on all vertices, so that the sole difference between the players is their positions in the network. Unless mentioned otherwise, $\mathcal{U} = \sqrt{x}$ in our simulations; on some simulations we chose $\mathcal{U} = x^p$ for various values for $0 < p < 1$.

Specifying the parameters of a model define a particular distribution on networks, and

We examined the properties of equilibria on 100 graphs sampled from each formation model. The distributions we studied in the preferential attachment model are $PA(50, 1)$

and $PA(50, k)$, for all integers k from 1 to 5, where $PA(n, k)$ denotes a preferential attachment model with 50 vertices and k new links being added per vertex. The distributions we studied in the Erdos-Renyi model are $ER(50, 4k/100)$, for all integers k from 1 to 5, where $ER(n, p)$ denotes a random graph with 50 vertices where the probability that an edge exists is p. Note that $ER(50, 4k/100)$ and $PA(n, k)$ have comparable number of edges in expectation. Unless specified otherwise, all the data presented below uses a utility function of $\sqrt{x}$ for all vertices.

On every network that was included in the sample, we ran random best response dynamics 20 times, for both PBS and NBS, each time starting from a random initial state and then updating unstable edges at random. In every network, this randomized algorithm converged to states that were all within a radius of 0.01 in l_∞ norm. Further, this radius had a decreasing trend whenever ϵ was reduced. This it seems plausible that on all these networks, every run of the algorithm converges to a small neighborhood of a unique equilibrium. This event occurred for both the PBS and NBS concepts. Since the approximate PBS equilibria computed are almost the same, we shall subsequently only consider one PBS equilibrium and one NBS equilibrium when analysing the wealth distribution on the vertices.

4.2 Correlation Between Degree and Wealth

Echoing earlier results found in a rather different (non-bargaining) model [13], we found that in both formation models there is a very high correlation between vertex degree and wealth at equilibrium — on average (over networks), correlations in excess of 0.95. Given such high correlations, it is natural to attempt to model the wealth of each vertex in a given network by a linear function of its degree — that is, in a given network we approximate the equilibrium wealth w_v of a vertex v of degree d_v by $\alpha d_v - \beta$, and minimize the mean squared error (MSE)

$$\frac{1}{n}\sum_{v\in V}(w_v - (\alpha d_v - \beta))^2$$

where $n = |V|$. We find that such fits are indeed quite accurate (low MSE). We do note that the correlation is generally higher, and the MSE generally lower, in NBS equilibrium compared to PBS equilibrium, so linear functions of degree are better models of wealth in NBS than PBS.

5 that minimize the squared error in each graph. The tables also report

Note that since the sum of the wealth of all vertices is equal to the number of edges m, we have

$$\sum_{v\in V}\alpha d_v - \beta \approx m \Rightarrow \alpha(2m) - \beta n \approx m \Rightarrow \beta \approx \frac{m}{n}(2\alpha - 1)$$

Thus, $\tilde{\beta} = \frac{m}{n}(2\alpha - 1)$ is an estimate of β that is almost accurate when the mean squared error is as small as we have found. $\tilde{\beta}$ is positive if $\alpha > 0.5$, and negative if $\alpha < 0.5$. In essence, therefore, the wealth distribution on the vertices of a specific network is succinctly really expressed by just a single real value – the degree coefficient α. Note that α itself is a function of the network, and as we shall see, is also dependent on the equilibrium concept.

4.3 Regression Coefficients

For each graph, we have two coefficient values α, one for the PBS equilibrium and the other for the NBS equilibrium.

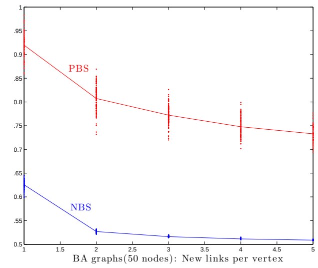

Figure 1: PBS and NBS regression coefficients versus edge density in PA

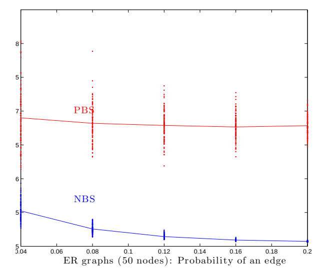

Figure 2: PBS and NBS regression coefficients versus edge density in ER

Thus each distribution of graphs gives two distributions of α, one for PBS equilibrium and the other for NBS equilibrium. A standard t-test reveals that for both network formation models, these distributions of coefficients have rather means at a high level of statistical significance. On all networks, the coefficient in the PBS equilibrium was higher than that in the NBS equilibrium, and both numbers were greater than 0.5.

Figure 1 shows the values of α for PBS and NBS equilibria in the 5 different distributions, which vary in their edge density, from the preferential attachment model. The horizontal axis shows the number of new links added per vertex in the random generation, while the vertical axis shows the regression coefficients from 100 trials. Figure 2 shows the analogous plot for distributions from the Erdos-Renyi model, with the horizontal axis representing the edge probability. The plots clearly demonstrate that the dependence of wealth on degree is stronger in PBS than in NBS. In preferential attachment networks, both equilibrium concepts seem to have diminished coefficients with increased edge density, but this effect is absent or muted in Erdos-Renyi.

4.4 Division of Wealth on Edges

So far we have examined the total wealth of *players* in equilibrium states; we can also examine how the wealth on individual edges is divided at equilibrium. Figures 3 and 4 show histograms of the division of wealth on edges in the distribution $PA(50, 2)$ for PBS and NBS equilibrium, respectively. The horizontal axis shows the amount of the smaller share of an edge, and the vertical axis shows the number of edges whose smaller share is within the given range. The number of edges is summed over 100 graphs from each distribution. In NBS, the wealth on most edges are split quite

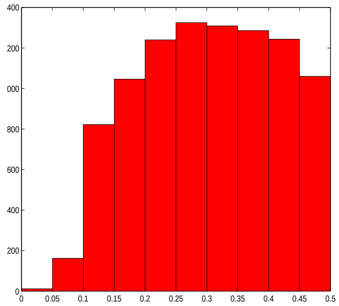

Figure 3: A histogram of PBS for division on each edge, for $PA(50,2)$

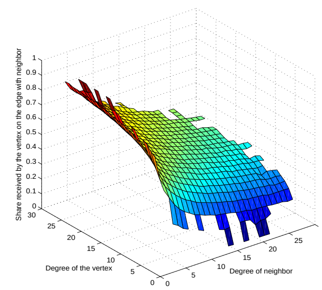

Figure 5: Edge division versus edge endpoint degrees in PBS for $PA(50,3)$

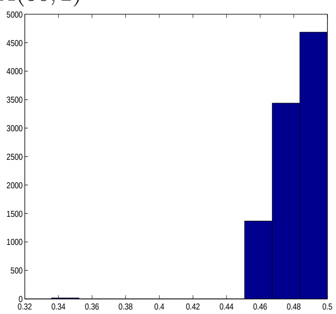

Figure 4: A histogram of NBS for division on each edge, for $PA(50,2)$

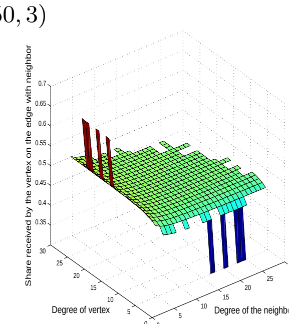

Figure 6: Edge division versus edge endpoint degrees in NBS for $PA(50,3)$

evenly, while in PBS, the split is heavily skewed, yet another indication that network structure plays a greater role in PBS equilibrium than in NBS equilibrium.

A more refined view of this phenomenon is provided in Figures 5 and Figure 6. Here we show the average edge divisions as a function of the degrees of the two endpoints d_1 and d_2 for the two equilibrium notions in the distribution $PA(50,3)$. We note that while both surfaces are smooth and have similar trends, the slope of the surface is much gentler in NBS equilibrium than in PBS equilibrium, demonstrating that even neighbors with rather different degrees tend to split deals approximately evenly at NBS equilibrium.

4.5 Other Utility Functions

All experiments described so far examined the utility function $\sqrt{x}$. We also performed experiments examining equilibria varied with a change of utility functions. We examined $\mathcal{U} = x^p$ for $p = i/10\ \forall 1 \leq i \leq 10$ on each graph distribution, and found that for each of them, the correlation of wealth and degree was still very high and linear fits still provide excellent approximations. We know theoretically that for $i = 10$, that is, $\mathcal{U}(x) = x$, we have $\alpha = 0.5$. Figure 7 illustrates how the degree coefficient for PBS equilibrium decreases smoothly in $PA(50,3)$, from an average value of almost 1 to 0.5, as p goes from 0.1 to 1.0, while that for NBS equilibrium starts barely above 0.5 and also goes down to 0.5, albeit with a far gentler slope. Figure 8 shows the same plot for the distribution $ER(50,0.12)$, which, in expectation, has approximately the same number of edges as in $PA(50,3)$. Thus, again viewing a higher value of α as a higher variance in bargaining power and thus greater effects of network structure, with $\alpha = 0.5$ implying the absence of network effect, we conclude from the figures that network effects gradually diminish when the utility function approaches linearity.

5. UNIQUENESS OF EQUILIBRIUM

The simulations suggested that equilibrium may be unique in the networks we ran our simulations on, or at least that the random best response dynamics converges to a unique one. In this section we address the question of whether equilibrium is unique. Unfortunately, the answer to this question is no.

In this section, we focus on regular graphs with unit wealth on all edges, and the same utility function $\mathcal{U}$, which is increasing, continuous and concave, for all vertices. This class has one state that is both PBS and NBS equilibrium: the state where the value on every edge is divided into two equal parts. We investigate if there is any other equilibrium. We show that one can choose $\mathcal{U}$ such that there are multiple equilibria, both for PBS and for NBS.

However, we give a simple condition of the update process that will ensure uniqueness in this class of networks. We show that many natural concave utility functions such as x^p, $0 < p < 1$ satisfy this condition in both PBS and NBS concepts, and thus ensure unique PBS and NBS equilibrium, respectively.

5.1 PBS Equilibrium is Not Unique

Consider any d-regular bipartite graph ($d \geq 2$) with edges of unit wealth, and every player with the same utility function $\mathcal{U}$, which is defined as $\mathcal{U}(x) = 100x$ for $0 \leq x \leq 0.01$, and $\mathcal{U}(x) = \log x - \log 0.01 + 1$ for $x > 0.01$. Then, $\mathcal{U}(x) = 0$, $\mathcal{U}$ is differentiable, and $\mathcal{U}'$ is a decreasing positive function, so $\mathcal{U}$ is concave and increasing. This bargaining network has uncountably many PBS equilibria.

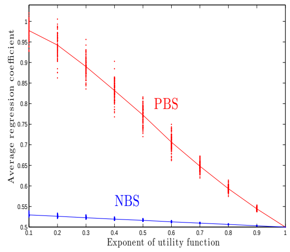

Figure 7: Regression coefficient against exponent of utility function, in $PA(50,3)$

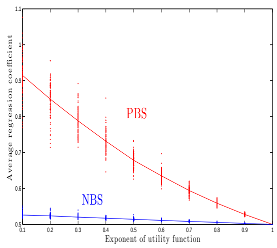

Figure 8: Regression coefficient against exponent of utility function, in $ER(50,12)$

Let the vertices in the bipartite graph be $X \cup Y$, where X and Y are independent sets. Consider a state s where on every edge, the endpoint in X receives $1/2 - \epsilon$, while the endpoint in Y receives $1/2 + \epsilon$, for any $0 < \epsilon < 0.49$. Then, for any edge, the outside options of its two endpoints are $(d-1)(1/2-\epsilon)$ and $(d-1)(1/2+\epsilon)$ respectively, and s is a PBS equilibrium if and only if $\mathcal{U}(d(1/2-\epsilon)) - \mathcal{U}((d-1)(1/2-\epsilon)) = \mathcal{U}(d(1/2+\epsilon)) - \mathcal{U}((d-1)(1/2+\epsilon))$. It is easy to check that both sides evaluate to $\log d/(d-1)$. Since the result holds for any $0 < \epsilon < 0.49$, we have a continuum of PBS equilibria.

5.2 NBS Equilibrium is Not Unique

We again consider d-regular bipartite graph ($d \geq 2$) with edges of unit wealth, and every player with the same utility function $\mathcal{U}$. However, we need to choose $\mathcal{U}$ more carefully, and we shall only get multiple NBS equilibrium, instead of uncountably many.

Let the vertices in the bipartite graph be $X \cup Y$, where X and Y are independent sets. Consider a state s where on every edge, the endpoint in X receives $1/4$, while the endpoint in Y receives $3/4$. Then, for any edge, the outside options of its two endpoints are $(d-1)/4$ and $3(d-1)/4$ respectively, and s is an NBS equilibrium if and only if

$$\frac{\mathcal{U}(d/4) - \mathcal{U}((d-1)/4)}{\mathcal{U}'(d/4)} = \frac{\mathcal{U}(3d/4) - \mathcal{U}(3(d-1)/4)}{\mathcal{U}'(3d/4)}$$

Note that $3(d-1)/4 > d/4$ when $d \geq 2$, so the above equation is easy to satisfy. We can define an increasing, continuous and concave function $\mathcal{U}$ such that $\mathcal{U}'(d/4) = 1$, $\mathcal{U}'(3d/4) = 1/2$ and $(\mathcal{U}(d/4) - \mathcal{U}((d-1)/4)) = 2(\mathcal{U}(3d/4) - \mathcal{U}(3(d-1)/4))$. In particular, we can choose $\mathcal{U}$ such that $\mathcal{U} = 8x \; \forall x \leq (d-1)/4$, so that $\mathcal{U}((d-1)/4) = 2(d-1)$, and then decrease the slope gradually such that $\mathcal{U}(d/4) = \mathcal{U}((d-1)/4) + 1 = 2d - 1$ and $\mathcal{U}'(d/4) = 1$, and then $\mathcal{U}'(3(d-1)/4) = 1$, and finally $\mathcal{U}(3d/4) = \mathcal{U}(3(d-1)/4) + 1/2$ and $\mathcal{U}'(3d/4) = 1/2$.

5.3 Uniqueness in Regular Graphs

The following condition, which is dependent on the utility function $\mathcal{U}$ of the players as well as the concept we are considering, PBS or NBS, is sufficient to ensure uniqueness on regular graphs with unit wealth on edges and identical utility functions for the players. We call it a rallying condition, because it allows the player with less outside options to bargain a greater share than what the disbalance in the outside options suggest, even though she gets the smaller portion.

Condition 5.1 (Bargain Rallying Condition) *Let s be any state of a bargaining network and $e = (u, v)$ be any edge of wealth c. Without loss of generality, suppose $\alpha_s(u, e) \leq \alpha_s(v, e)$. Then, our condition is that $c/2 \geq y_s(u, e) > \frac{c\alpha_s(u,e)}{\alpha_s(u,e)+\alpha_s(v,e)}$.*

This condition is satisfied by common concave utility functions such as $\mathcal{U}(x) = x^p$ where $0 < p < 1$, and also $\mathcal{U}(x) = f + \log(a + bx)$ where $a > 0,\ b > 0$. A detailed discussion is deferred to the full version.

We shall now present and prove our main uniqueness result.

Theorem 3 *If Condition 5.1 is satisfied in the PBS (or NBS) concept on a regular graph with unit wealth on every edge and identical utility function for all players, which is increasing, continuous and concave, then there is a unique PBS (or NBS, respectively) equilibrium, where the wealth on every edge is shared equally between its endpoints.*

PROOF. Proof by contradiction. Clearly, the state, where wealth on every edge is equally divided, is an equilibrium. Suppose there exists another equilibrium s. There exists an edge where $x_s(u, e) < 1/2$ or $x_s(v, e) < 1/2$. Consider the edge $e = (u, v)$ which has the most lopsided division, that is, find e such that $\min(x_s(u, e), x_s(v, e))$ is minimized. Without loss of generality, suppose that u gets the smaller share $1/2 - \epsilon$ on e, for some $\epsilon > 0$. Since Condition 5.1 is satisfied, so $\alpha_s(u, e) \leq \alpha_s(v, e)$. Note that u gets at least $1/2 - \epsilon$ from each edge incident on it, and so $\alpha_s(u, e) \geq (d-1)(1/2-\epsilon)$. Similarly, v gets at most $1/2 + \epsilon$ from each incident edge, so $\alpha_s(v, e) \leq (d-1)(1/2+\epsilon)$. Thus $\alpha_s(v, e)/\alpha_s(u, e) \leq \frac{1/2+\epsilon}{1/2-\epsilon}$. Now using Condition 5.1, we get that $y_s(u, e) > \frac{1}{1+\frac{1/2+\epsilon}{1/2-\epsilon}} = 1/2 - \epsilon$ which is a contradiction to our assumption that s is an equilibrium. □

6. GENERALIZED UTILITY FUNCTIONS

In all the results in this paper, as well as the definitions in the Preliminary section, utility function was assumed to be a function of the sum of the wealth received by the player on all edges incident upon it, that is the function was of the form $\mathcal{U}(x_1 + x_2 \ldots + x_d)$, where $x_1, x_2 \ldots x_d$ are the different shares she receives from her d edges. We can, instead, define a more general concept of a utility function, one that is multi-dimensional, and is an arbitrary function of the values of the wealth she receives from her edges, that is, of the form

$\mathcal{U}(x_1, x_2 \ldots x_d)$. The concepts of PBS and NBS equilibrium can be easily extended to this setting, by simply redefining differential utility, which is still the additional utility a player receives from the deal, given what she is getting from her other deals. In particular, her differential utility from the first incident edge is $\mathcal{U}(x_1, x_2, x_3 \ldots x_d) - \mathcal{U}(0, x_2, x_3 \ldots x_d)$. Note that the outside option for this first edge should also be redefined as a sequence of d values $(0, x_2, x_3 \ldots x_d)$ where the value corresponding to the first edge is zero, while the rest reflects the wealth the player receives from her other $(d-1)$ deals.

We think it is worth noting that even under this generalized model, the existence results of [5] hold, that is,

- PBS equilibrium exists on all networks if the utility functions are increasing and continuous.
- NBS equilibrium exists on all networks if the utility functions are increasing, continuous and concave.

These results follow from the fact that analogous versions of Lemmas 2.1 and 2.2 hold in this model too, using the fact that an increasing, continuous and/or concave multi-dimensional function has the same property along each dimension.

Further, it is easy to verify that the bargaining monotonicity condition is satisfied by concave utility functions in the PBS concept. Again, the condition needs to be modified to take care of the fact that outside options are now a sequence: we say that $\alpha_s(u,e) \geq \alpha_{s'}(u,e)$ if the former dominates the latter in every dimension. This condition is sufficient to show that the same FPTAS algorithm works even in this generalized model.

7. CONCLUSION

Our model is an addition to the extensive literature on network exchange theory, and it differs from previous models in that network effects are caused by the non-linearity of utility functions of the players. When players can enter into as many deals as their degree our model allows the possibility that not all edges will be split equally. Effects of network structure may be felt throughout the network when utility functions are concave.

The most prominent theoretical question left unanswered in our model is that of computing approximate or exact equilibria in general graphs.

Acknowledgements

We thank Jinsong Tan for insightful discussions about the formulation of our model, and Yishay Mansour for valuable technical suggestions.

8. REFERENCES

[1] A. L. Barabasi and R. Albert. Emergence of scaling in random networks. *Science*, 286:509–512, 1999.

[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 and M. Kearns. Bargaining solutions in a social network (short paper). In *Workshop on Internet and Network Economics (WINE)*, pages 548–555, 2008.

[6] K. S. Cook and T. Yamagishi. Power in exchange networks: A power-dependence formulation. *Social Networks*, 14:245–265, 1992.

[7] S. Currarini, M. O. Jackson, and P. Pin. An economic model of friendship: Homophily, minorities and segregation. *Econometrica (forthcoming)*, 2009.

[8] X. Deng and C. H. Papadimitriou. On the complexity of cooperative solution concepts. *Mathematics of Operations Research*, 19(2):257–266, 1994.

[9] P. Erdos and A. Renyi. On random graphs. *Publicationes Mathematicae*, 6:290–297, 1959.

[10] N. Friedkin. An expected value model of social power: Predictions for selected exchange networks. *Social Networks*, 14:213–230, 1992.

[11] M. O. Jackson. Average distance, diameter, and clustering in social networks with homophily. In *Workshop on Internet and Network Economics (WINE)*, pages 4–11, 2008.

[12] S. Judd and M. Kearns. Behavioral experiments in networked trade. In *ACM Conference on Electronic Commerce (EC)*, pages 150–159, 2008.

[13] S. M. Kakade, M. J. Kearns, L. E. Ortiz, R. Pemantle, and S. Suri. Economic properties of social networks. In *Neural Information Processing Systems Conference (NIPS)*, 2004.

[14] 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.

[15] J. Kleinberg and E. Tardos. Balanced outcomes in social exchange networks. In *Symposium on Thoery of Computing (STOC)*, pages 295–304, 2008.

[16] 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.

[17] 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.