The Economics of Social Networks

Sham M. Kakade **Michael Kearns** **Luis E. Ortiz** **Robin Pemantle** **Siddharth Suri**

University of Pennsylvania
Philadelphia, PA 19104

Abstract

In an attempt to explain a number of apparently universal statistical properties of natural social and economic networks, many authors have recently proposed a rich family of probabilistic generative models for social network formation. We have recently introduced a graph-theoretic representation for classical mathematical economic models. This paper studies the marriage of these two lines of inquiry: we examine the network effects that social network models (such as preferential attachment) introduce into classical economic settings. Our findings are a mixture of theoretical results and large-scale simulations that have been made possible only by algorithmic advances of the last two years.

1 Introduction

In this paper, we report on theory and experiments in a setting that blends classical economic exchange models with more recent proposals for how "natural" social and economic networks form. The glue that permits this interesting marriage is the recently-introduced framework of *graphical economics*.

Within economics, there is a long history of research on mathematical models for exchange markets, and the existence and properties of their equilibria. The work of Arrow and Debreu [1954], who established equilibrium existence in a very general commodities exchange model, was certainly one of the high points of this continuing line of inquiry. The origins of the field go back at least to Fisher [1891].

While there has been relatively recent interest in network models for interaction in economics, it was only quite recently that a network or graph-theoretic model that contains and generalizes the classical Arrow-Debreu and Fisher models was introduced (Kakade et al. [2004]). In this model, the edges in a network over individual consumers (for example) represent those pairs of consumers that can engage in direct trade. As such, the model captures the many real-world settings that can give rise to limitations on the trading partners of individuals (regulatory restrictions, social connections, embargoes, and so on). In addition, variations in the price of a good can arise due to the topology of the network: certain individuals may be relatively favored or cursed by their position in the graph.

In a parallel development over the last decade or so, there has been an explosion of interest in what is broadly called *social network theory* — the study of apparently "universal" properties of natural networks (such as small diameter, local clustering of edges, and heavy-tailed distribution of degree), and statistical generative models that explain such properties. While there are certainly instances of "economic" properties or problems being examined in such generative models (see Jackson [2003] for a good review), the assumptions of individual rationality in these works are usually either non-existent, or quite weak compared to the Arrow-Debreu or Fisher models.

The introduction of graphical economics permits the examination of the classical economic exchange models in the modern light of social network theory, and this is the topic of this paper. Of central interest are two broad categories of question:

- To what extent can the combination of classical economic models and recent generative models for networks "explain" longstanding economic phenomena? A sample question of this kind is: Does the *preferential attachment* model of network formation (Barabasi and Albert [1999]), combined with the Fisher model of economic equilibrium, predict the heavy-tailed distribution of wealth first observed by Pareto?

- How do the properties of economic equilibrium vary with the statistical properties of the network formation process? For instance, while it is clear that significant price variation can occur in very sparse networks, and that there can be no such variation at equilibrium in the complete graph, our particular setting lets us carefully study the transition from significant to no variation as a function of edge density and other network parameters.

This paper describes our initial theoretical and experimental investigations into such questions. More specifically, we establish the following results:

- The tails of the wealth distribution at economic equilibrium in preferential attachment networks obeys a power law, but one that is different from the degree distribution. These tails are rapidly diminished as we increase connectivity in various ways.
- Price variation can be great in such networks, scaling as a power of network size, but is rapidly diminished by less preferential attachment.

Many of our results are based on a powerful new *local approximation* method for global equilibrium prices: we show that in the preferential attachment model, prices computed from only local regions of a network yield strikingly good estimates of the global prices. We exploit this method computationally and theoretically.

2 Economic and Network Models

In this section, we define our economic models, and the generative model for social networks that we consider.

2.1 The Graphical Linear Fisher Model

We first describe the standard *Fisher model*, which consists of a set of *buyers* and a set of *goods*. We assume that there are g_j units of good j in the market, and that each good j is be sold at some price p_j. Each buyer i has a cash *endowment* e_i, to be used to purchase goods in a manner that maximizes the buyer's utility. In this paper we make the well-studied assumption that the utility function of each buyer is *linear* in the amount of goods consumed (see Gale [1960]), and leave the more general case to future research. Let $u_{ij} \geq 0$ denote the utility derived by i on obtaining a single unit of good j. If i consumes x_{ij} amount of good j, then the utility i derives is $\sum_j u_{ij}x_{ij}$.

A set of *prices* $\{p_j\}_j$ and *consumption plans* $\{x_{ij}\}_{i,j}$ constitutes an *equilibrium* if the following two conditions hold:

1. The market *clears*, *i.e.* supply equals demand. More formally, for each j, $\sum_i x_{ij} = g_j$
2. The consumption plan $\{x_{ij}\}$ is optimal for each buyer i. By this we mean that the consumption plan maximizes the linear utility function of i, subject to the constraint that the total cost of the goods purchased by i is not more than the endowment e_i.

It turns out that such an equilibrium always exists if each good j has a buyer which derives nonzero utility for good j — that is, $u_{ij} > 0$ for some i (see Gale [1960]). Furthermore, the equilibrium prices are unique.

We now consider the *graphical Fisher model*, so named because of the introduction of a graph-theoretic or network structure to exchange. In the basic Fisher model, we implicitly assumed that all goods were available in a centralized exchange, and all buyers had equal access to these goods. In the graphical Fisher model, we would like to capture the fact that each good may have multiple vendors or *sellers*, and that individual buyers may have access only to some, but not all, of these sellers. There are innumerable settings where such asymmetries arise. Examples include the fact that consumers generally purchase their groceries from local markets, that social connections play a major role in business transactions, and that securities regulations prevent certain pairs of parties from engaging in stock trades.

Without loss of generality, we assume each that seller j sells only one of the available goods. (Each good may have multiple competing sellers.) Let G be a bipartite graph, where buyers and sellers are represented as vertices, and all edges are between a buyer-seller pair. The semantics of the graph are as follows: if there is an edge from buyer i to seller j, then buyer i is permitted to purchase from seller j. Note that if buyer i is connected to two sellers of the same good, he will always choose to purchase from the cheaper source, since his utility is identical for both sellers (they sell the same good).

The graphical Fisher model is a special case of a more general and recently introduced framework (Kakade et al. [2004]). One of the most interesting features of this model is the fact that at equilibrium, significant price variations can appear solely due to structural properties of the underlying network. The current paper is the first to examine price variation, wealth distribution, and other economic metrics in recent statistical models for network generation.

It is easy to see that if G is the complete bipartite graph between buyers and sellers, and there is a single seller for each good j, we recover the standard Fisher model. Alternatively, we can encode a graphical Fisher economy in the standard model as follows: for each seller s of good j, introduce the virtual good (j, s). For buyer i, if u_{ij} is the utility for the (original) good j, then let u_{ij} be the utility that i has for (j, s) if s is a neighbor of i in G; and let the utility for all other virtual goods be 0. In this manner the utilities for virtual goods encode both the original utilities and the structure of G.

Thus, the mere introduction of the graphical Fisher model is no advance over the classical model. The novelty of the results described here lies in the examination of how the *structure* of G — in particular, the statistical structure produced by now-standard generative models from social network theory — influences properties of economic equilibrium. We now describe such a generative model.

2.2 Preferential Attachment Networks

For simplicity, in the sequel we will without loss of generality consider economies in which the numbers of buyers and sellers are equal. We will also restrict attention to the case in which all sellers sell the *same* good. We note that from a mathematical and computational standpoint, this restriction is rather weak: when considered in the graphical setting, it already contains the setting of multiple goods with binary utility values. (See remarks in the last section about encoding the graphical setting in the classical Fisher model via the introduction of virtual goods.)

The simplest generative model for the bipartite graph G might be the *random graph*, in which each edge between a buyer i and a seller j is included independently with probability p. This is simply the bipartite version of the classical Erdos-Renyi model (Bollobas [2001]).

Many researchers have sought more realistic models of social network formation, in order to explain observed phenomena such as heavy-tailed degree distributions. We now describe a slight variant of the *preferential attachment* model (Mitzenmacher [2003]) for the case of a bipartite graph. We start with a graph in which one buyer is connected to one seller. At each *time step*, we add one buyer and one seller as follows. With probability α, the buyer is connected to a seller in the existing graph uniformly at random; and with probability $1 - \alpha$, the buyer is connected to a seller chosen *in proportion to the degree* of the seller (preferential attachment). Simultaneously, a seller is attached in a symmetric manner: with probability α the seller is connected to a buyer chosen uniformly at random, and with probability $1 - \alpha$ the seller is connected under preferential attachment.

The parameter α in this model allows us to move between a pure preferential attachment model ($\alpha = 0$), and a model closer to classical random graph theory ($\alpha = 1$), in which new parties are connected to random extant parties. We note that the latter still does not exactly produce the Erdos-Renyi model due to the incremental nature of the network generation: early buyers and sellers are still more likely to have higher degree. However, this bias is rather weak, as we shall see.

Note that the above model always produces trees, since the degree of a new party is always 1 upon its introduction to the graph. We thus will also consider a variant of this model in which at each time step, a new seller is still attached to exactly one extant buyer, while each new buyer is connected to $\nu > 1$ extant sellers. The procedure for edge selection is as outlined above, with the modification that the ν new edges of the buyer are added without replacement — meaning that we resample so that each buyer gets attached to exactly ν distinct sellers.

The main purpose of the introduction of ν is to have a model capable of generating highly cyclical (non-tree) networks, while having just a single parameter that can "tune" the asymmetry between the (number of) opportunities for buyers and sellers. However, there are less mathematical motivations as well: it is natural to imagine that new sellers of the good arise only upon obtaining their first customer, but that new buyers arrive already aware of several alternative sellers.

In the sequel, we shall refer to the generative model just described as the *bipartite* (α, ν)*-model*. We will use n to denote the number of buyers and the number of sellers, so the network has $2n$ vertices.

2.3 An Illustrative Example

As an illustration of the marriage of economic and social network models we have introduced, Figure 2.3 shows a sample graph generated by the bipartite ($\alpha = 0, \nu = 2$)-model with $n = 15$. Buyers and sellers are labeled by 'B' or 'S' respectively, followed by an index indicating the time step at which they were introduced to the network. Each seller is labeled with the price they charge at equilibrium. Note that in this example, there is non-trivial price variation, with the most fortunate sellers charging 2.00 for their unit of the good, and the least fortunate 0.33. Also note that while there appears to be a correlation between seller degree and price, it is far from a deterministic relation, a topic we shall examine.

The solid edges in the figure show the *exchange subgraph* — those pairs of buyers and sellers who actually exchange currency and goods at equilibrium. Note the

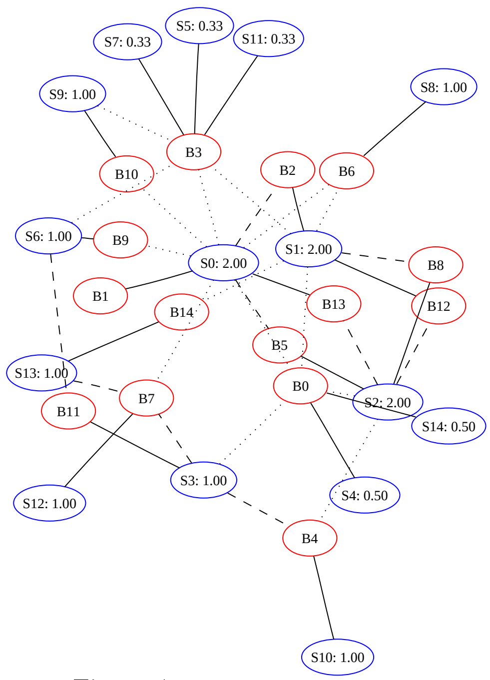

Figure 1: See text for description.

sparseness of this graph compared to the overall graph. The dotted edges are edges of the network that are unused at equilibrium because they represent inferior prices for the buyers, while the dashed edges are edges of the network that have competitive prices, but are unused at equilibrium due to the specific consumption plan required for market clearance.

3 Statistics of the Network

In this and the following section we summarize our theoretical findings. We begin with results establishing purely statistical (non-economic) properties of the bipartite (α, ν)-model, and then apply these to establish economic properties of the model. We note that while the statistical results are reminiscent of the recent literature (Mitzenmacher [2003]), we need to establish their dependence on α and ν, as well as deal with the bipartite structure of the graph.

We first establish that the degree distribution of the sellers obeys a power law. (In the sequel, we only examine properties of the sellers; similar statements hold for buyers.) Let $Y(j,n)$ denote the degree of the j^{th} seller (in order of arrival) at time $n \geq j$ (that is, after n buyers and n sellers have been added to the network). The following lemma characterizes the behavior of the random variable $Y(j,n)$ asymptotically. Let

$$\beta = (1-\alpha)\nu/(1+\nu)$$

Lemma 1 *In the bipartite (α, ν)-model, $Y(j,n)$ tends to $(1+\frac{\alpha\nu}{\beta})(n/j)^{\beta}$ for sufficiently large n.*

Proof: (Sketch) Establishing this lemma rigorously is beyond the scope of this paper, so we only provide a non-rigorous argument with respect to the means $y_{j,n} = \mathbf{E}[Y(j,n)]$. A long version of this paper will have the complete proof.

The total number of edges after time n is $(1+\nu)n$. From time n to $n+1$, each of the ν additional edges is attached to seller j with probability $(1-\alpha)y_{j,n}/((1+\nu)n) + \alpha/n$. By linearity of expectation, we can sum over the ν edges without worrying about the negative dependence arising from sampling without replacement, which implies

$$y_{j,n+1} = y_{j,n}\left(1 + \frac{(1-\alpha)\nu}{(1+\nu)n}\right) + \alpha\frac{\nu}{n}$$

In the forthcoming version, we solve this formula exactly, but here we treat the system as the differential equation (as in Mitzenmacher [2003])

$$\frac{dy_{j,n}}{dn} = \frac{\beta y_{j,n} + \alpha\nu}{n}$$

where $\frac{dy_{j,n}}{dn}$ is used for $y_{j,n+1} - y_{j,n}$. Solving this with the boundary condition $y_{j,j} = 1$ leads to the result. □

We may now translate this result into a power law for the distribution of seller degrees.

Theorem 2 *In the bipartite (α, ν)-model, for $x = o(n^{1/\beta})$, the proportion of sellers at time n whose degree exceeds x is $\Theta(x^{-1/\beta})$.*

Proof: (Sketch) Fix n, and consider the proportion of sellers whose degree exceeds some value x. That is, consider the j that solves $y_{j,n} = x$. Hence, $j = \Theta(nx^{-1/\beta})$. Dividing by n provides the proportion of sellers, which is the desired result. In the long version, we show that the approximation of only considering means is sufficient. □

For the simplest case of $\alpha = 0$ and $\nu = 1$, the tail of this cumulative degree distribution is just $\Theta(x^{-2})$. More generally, as α approaches 1 (towards unbiased selection, and away from preferential attachment), the exponent blows up, and the tails of the distribution become lighter. At $\alpha = 1$, we actually have exponential rather than power law decay.

4 Economics of the Network

We next present a rather intuitive "monotonicity" lemma, which states that if the supply of goods in a classical Fisher economy is decreased, or the cash endowments are increased, then the equilibrium prices do not decrease. We will then apply this lemma in the graphical Fisher model, along with the results of the previous section, to obtain the distribution of seller wealth.

Lemma 3 *(Monotonicity) Let E and E' be two Fisher economies with the same number of buyers and sellers and identical linear utility functions. If for all goods j and buyers i, we have $g'_j \leq g_j$ and $e'_j \geq e_j$ (where the primes denote quantities for economy E'), then the equilibrium prices satisfy $p'_j \geq p_j$ for all j.*

Proof: To prove this, we use properties of a recent algorithm for computing equilibria in the linear Fisher model (see Devanur et al. [2002]), which we now describe. Define the "bang per buck" for buyer i consuming good j at price $\tilde{p}_j$ as $u_{ij}/\tilde{p}_j$. Clearly, it is only optimal for buyer i to purchases those goods which have maximal bang per buck.

The algorithm is an iterative scheme in which prices $\{\tilde{p}_j\}$ are increased at every iteration, until an equilibrium is reached. Importantly, the algorithm can be initialized to any prices which obey the following property, which is referred to as the "Invariant" in Devanur et al. [2002][1]. We say that the Invariant holds at prices $\{\tilde{p}_j\}$ if the buyers have enough money to purchase all the goods in the market, while only purchasing goods which maximize their bang per buck (though the buyers may have left over money after this purchase). Essentially, the Invariant holds at some prices if the buyers can clear the market while purchasing optimally at these prices.

Let us now use this algorithm to compute the equilibrium prices in economy E'. It suffices to show that we can initialize this algorithm to the equilibrium prices of E, $\{p_j\}$, since the algorithm only increases the prices. To show that such an initialization is sound, we only need to show that the prices $\{p_j\}$ satisfy the Invariant in E'.

To show this, first note that since these prices are an equilibrium in E, then the buyers can use their money endowments of $\{e_j\}$ to clear an amount of goods $\{g_j\}$, while only purchasing goods which maximize their bang per buck. Hence, by assumption, the buyers in E' can use larger money endowments of $\{e'_j\}$ to clear a smaller amount of goods $\{g'_j\}$, while only purchasing goods which maximize their bang per buck (since the utility functions in E and E' are identical). □

This lemma immediately implies a scheme in which we can find upper and lower bounds on the equilibrium prices using only *local* computations. First, note that any subset V' of buyers and sellers defines a natural *induced economy*, where the induced graph G' consists of all edges between buyers and sellers in V' that are also in G. We say that G' has a *buyer (respectively, seller) frontier* if on every (simple) path in G from a node in V' to a node outside of V', the last node in V' on this path is a buyer (respectively, seller).

Corollary 4 *(Frontier Bounds) If V' has a subgraph G' with a seller (or buyer) frontier, then the equilibrium price of any good j in the induced economy on V' is a lower bound (or, respectively, an upper bound) on the equilibrium price of j in G.*

Proof: Let us prove the lower bound for the seller frontier case. Consider setting the cash e_i of all buyers i not in G' to 0. By the previous lemma, the equilibrium prices in this modified economy E' is a lower bound on the equilibrium prices for the economy with graph G. Note that all sellers in G' have no demand from any buyers outside of G', and, by definition of G', all buyers in G' purchase goods only from sellers in G'. So the equilibrium prices in the induced economy on G' are identical to their respective prices in E'. A symmetrical argument proves the upper bound case. □

This corollary has both statistical implications and computational implications. We now investigate the statistical implications in the (α, ν)-model, and examine the computational ones in Section 5.

Corollary 4 implies a simple wealth upper bound: the wealth of any seller j is bounded by its degree d. (By the *wealth* of a seller, we mean the price at which that seller sells their good. This terminology is justified by the fact that at equilibrium, this is exactly the income the seller receives after selling their one unit of good, since the market clears.) Although the same upper bound can be seen from first principles, it is instructive to apply Corollary 4. Let G' be the immediate neighborhood of j (which is j and its d buyers); then the equilibrium price in G' is just d, since all d buyers are forced to buy from seller j. This provides an upper bound since G' has a buyer frontier. Since the degree distribution obeys a power law in the bipartite (α, ν)-model, we have an upper bound on the cumulative wealth distribution.

Corollary 5 *In the bipartite (α, ν)-model, for $w = o(n^{1/\beta})$, the proportion of sellers with wealth greater than w is $O(w^{-1/\beta})$.*

We do not yet have such a closed-form lower bound on the cumulative wealth distribution. However, as we shall see in Section 5, the wealth distributions seen in large simulation results do indeed show power-law behavior. Interestingly, this occurs despite the fact that degree is a *poor* predictor of *individual* seller wealth.

[1] Devanur et al. [2002] choose a particular initialization, but it is clear that the algorithm is sound for any choice of initial prices which obey the Invariant.

Another quantity of interest is what we might call price or wealth variation — the ratio of the wealth of the richest seller to the poorest seller. The following theorem addresses this.

Theorem 6 *In the bipartite (α, ν)-model, if $\alpha(\nu^2 + 1) < 1$, then the ratio of the maximum price to the minimum price at time n is $\Omega(n^{\frac{2-\alpha(\nu^2+1)}{1+\nu}})$.*

For the simplest case in which $\alpha = 0$ and $\nu = 1$, this lower bound is just $\Omega(n)$.

Proof: (Sketch) Using Lemma 3, it straightforward to show the following two bounds on the maximum and minimum price. Consider the first ν sellers (in order of time) and let m be the number of buyers that are *only* connected to these ν sellers. Hence, the total wealth of these sellers must be m, so one of the first ν sellers must have a price that is m/ν, which is a lower bound on the maximum price. Similarly, an upper bound on the minimum price is provided by the price the first buyer obtains for his purchases, p_b. Equivalently, we use a lower bound $1/p_b$, which is the amount of goods this buyer purchases. This lower bound is provided by those sellers which are *only* connected to the first buyer.

Let us now bound the total wealth of the first ν sellers. The degrees of these sellers at time $n/2$ are all $\Theta(n^{\beta})$, so when a buyer arrives at a time between $n/2$ and n, the probability of one of this buyer's connections links to exactly the first ν sellers is $\Theta(n^{\beta-1})$. Hence, the probability that all of this buyer's connections link to exactly the first ν sellers is $\Theta(n^{\nu\cdot(\beta-1)})$. Summing over the $n/2$ buyers shows that the total number of such buyers is, with high probability, $\Theta(n^{1+\nu\cdot(\beta-1)})$. Deleting from this list those buyers who are later linked by some seller removes a constant fraction of these (shown in the long version). Hence, the first ν sellers have at least $\Omega(x^{\nu\cdot(1-\beta)})$ total wealth, which implies that the richest seller must have at least this wealth (treating ν is a constant).

Using similar arguments as in the proof of Lemma 1, one can show the first buyer has degree $\Theta(n^{\frac{1-\alpha}{1+\nu}})$. A similar argument to above shows that the $1/p_b$, which is the number of sellers *only* connected to this buyer, is $\Omega(n^{\frac{1-\alpha}{1+\nu}})$. Combining the previous bounds leads to the result. □

This proof can be generalized to obtain bounds on ratio of the wealth contained among the top x percent of sellers versus the poorest x percent of buyers (which is more of a relevant quantity in large economies).

5 Experimental Findings

We now present a number of experimental findings. Our equilibrium computations are done using the algorithm of Devanur et al. [2002] (or via the application of this algorithm to local subgraphs), which is described in the proof of Lemma 3. We note that it was only the recent development of this algorithm and related ones that made possible the simulations described here (involving hundreds of buyers and sellers in highly cyclical graphs). However, even the speed of this algorithm limits our experiments to networks with $n = 250$ if we wish to run repeated trials to reduce variance. Many of our results suggest that the local approximation schemes discussed below in Section 5.2 may be far more effective.

All simulations were performed on networks generated according to the bipartite (α, ν)-model.

5.1 Wealth and Degree Distributions

The leftmost panel of Figure 2 shows empirical *cumulative* wealth and degree distributions on a loglog scale, averaged over 25 networks drawn according to the bipartite $(\alpha = 0.4, \nu = 1)$-model with $n = 250$. The cumulative degree distribution is shown as a dotted line, where the y-axis represents the fraction of the sellers with degree greater than or equal to d, and the degree d is plotted on the x-axis. Similarly, the solid curve plots the fraction of sellers with wealth greater than some value w, where the wealth w is shown on the x-axis. The thin sold line has our theoretically predicted slope of $\frac{-1}{\beta} = -3.33$, which shows that degree distribution is quite consistent with our expectations, at least in the tails.

Let us examine some of the properties of this plot. Clearly, the cumulative degree distribution is flat until the degree is 1, since all sellers have a degree that is at least 1. In contrast, the cumulative wealth distribution quickly drops below 1, since a seller's wealth is not lower bounded by 1. (For instance, if a seller's potential buyers have many choices, then the seller's wealth could be driven arbitrarily close to 0.)

Perhaps the most interesting finding is that the tail of the *wealth* distribution looks linear, *i.e.* it also exhibits power law behavior. Our theory provided an upper bound, which is precisely the cumulative degree distribution. We do not yet have a formal lower bound. This plot (and other experiments we have done) further confirm the robustness of the power law behavior in the tail, for $\alpha < 1$ and $\nu = 1$.

As discussed in the Introduction, Pareto's original observation was that the wealth distribution in societies obey a power law, which has been born out in many

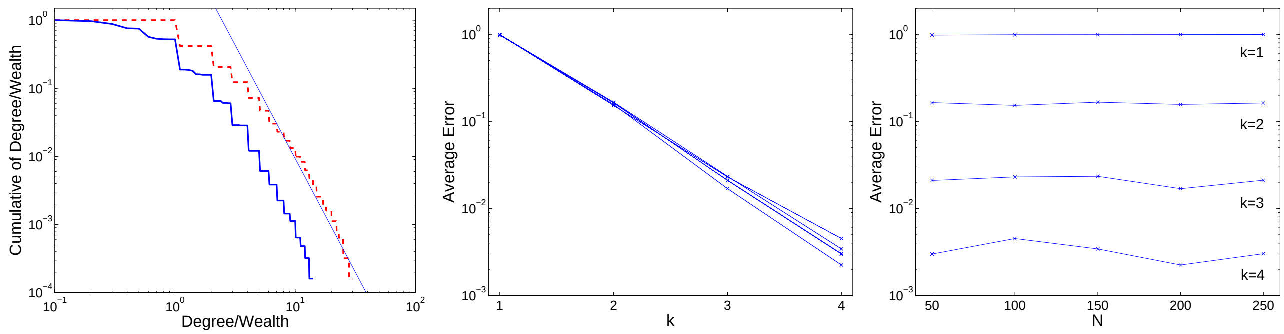

Figure 2: See text for descriptions.

studies on western economies. Since Pareto's original observation, there have been too many explanations of this phenomena to recount here. However, to our knowledge, all of these explanations are more *dynamic* in nature — a macroscopic dynamical system of wealth exchange leads to a power law tail, but without capturing the microscopic properties of individual rationality. Here we have power law wealth distribution arising from the combination of certain natural statistical properties of the network (generated by preferential attachment), and classical theories of economic equilibrium.

At this point it would be natural to conjecture that the wealth of a seller is essentially determined by its degree. We shall shortly see that the picture is more complicated, and that while degree is a rather poor predictor of individual seller wealth, more complex (but still local) properties are extremely accurate predictors.

5.2 Bounds via Local Computations

Recall that Corollary 5 suggests a scheme by which we can do only *local* computations to approximate the *global* equilibrium price for any seller. More precisely, for some seller j, consider the subgraph which contains all nodes that within distance k of j. In our bipartite setting, for k odd, this subgraph has a buyer frontier, and for k even, this subgraph has a seller frontier, since we start from a seller. Hence, the equilibrium computation on the odd k (respectively, even k) subgraph will provide an upper (respectively, lower) bound.

This provides an heuristic in which one can examine the equilibrium properties of small regions of the graph, without having to do expensive global equilibrium computations. The effectiveness of this heuristic will of course depend on how fast the upper and lower bounds tighten. In general, it is possible to create specific graphs in which these bounds are arbitrarily poor until k is large enough to encompass the entire graph. As we shall see, the performance of this heuristic is dramatically better in the bipartite (α, ν)-model.

The center panel in Figure 2 shows how rapidly the local equilibrium computations converge to the true global equilibrium prices as a function of k. On the x-axis is the value of k, and the y-axis shows the average difference (across all sellers) between the locally computed and global equilibrium prices on a log scale. Each line was created using different sized graphs (from $n = 250$ to $n = 50$, in increments of 50), and each line averages over 5 graphs. In these experiments, graphs were generated by the bipartite $(\alpha = 0, \nu = 1)$ model.

The linear nature of these plots establishes the fact that the error of the local approximations is decaying *exponentially* with increased k — indeed, by examining only neighborhoods of 3 steps from a seller in an economy of hundreds, we are already able to compute approximations to global equilibrium prices with errors in the second decimal place. The approximation $k = 1$ corresponds exactly to using seller degree as a proxy for price, and we can see that this performs rather poorly.

The right panel in Figure 2 shows a different view of the same data, and demonstrates how the average error scales with n, for each fixed value of k. It appears that for each value of k, the quality of approximation obtained is either independent of n, or has an imperceptibly mild dependence at the values of n for which we can compute global equilibria.

It turns out that $k = 5$ typically returned the exact equilibrium, even when $n = 250$. Furthermore, the diameter for $n = 250$ was often about 17, so the local graph is often considerably smaller than the global. Computationally, we found that the time required to do all 250 local computations for $k = 3$ was about 60% less than the global computation, and would result in presumably greater savings at much larger values of n.

5.3 Parameter Dependencies

We conclude with a brief examination of how wealth distribution and price variation depend on the param-

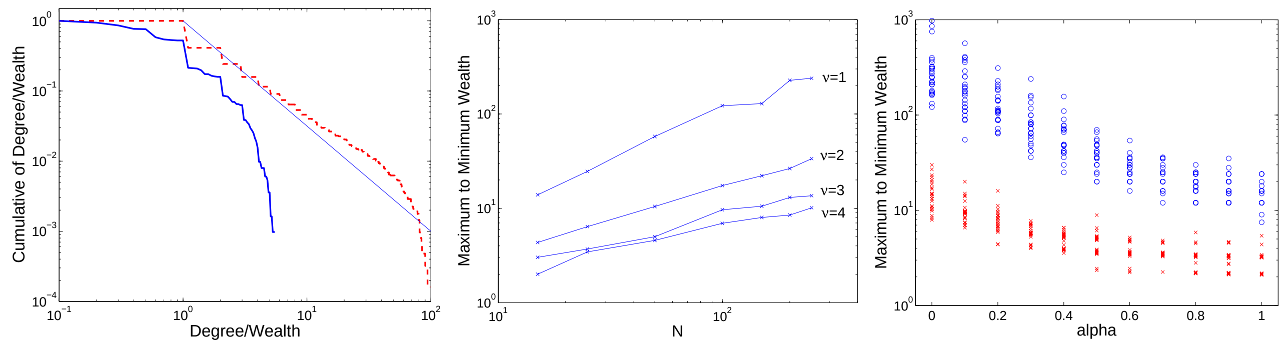

Figure 3: See text for descriptions.

eters of the bipartite (α, ν)-model.

The left panel of Figure 3 again shows *cumulative* wealth and degree distributions, as in Figure 2 . This time the values $\alpha = 0$ and $\nu = 2$ were used, and we see that the wealth distribution now deviates significantly from the degree distribution. Informally, we find that as ν increases, the tail of the wealth distribution deviates more sharply from the power law tail of the cumulative degree distribution. Lower bounds would help in establishing if the behavior is actually a power law behavior for $\nu > 1$.

We now turn to experimentally probing the lower bounds provided in Theorem 6. The center panel of Figure 3 shows the maximum to minimum wealth as function of n (averaged over 25 trials) on a loglog scale. Each line is for a fixed value of ν, and the values of ν range form 1 to 4 ($\alpha = 0$).

Recall from theorem 6, our lower bound on the ratio is $\Omega(n^{\frac{2}{1+\nu}})$ (using $\alpha = 0$). We conjecture that this lower bound is tight. If this is so, then the slopes of lines (in the loglog plot) should be $\frac{2}{1+\nu}$, which would be $(1, 0.67, 0.5, 0.4)$. The estimated slopes are somewhat close: $(1.02, 0.71, 0.57, 0.53)$.

The overall message is that for small values of ν, price variation increases rapidly (both theoretically and experimentally) with the economy size n in preferential attachment.

The right panel of Figure 3 is a scatter plot of α vs. the maximum to minimum wealth in a graph (where $n = 250$) . Here, each point represents the maximum to minimum price ratio in a specific network generated by our model. The circles are for economies generated with $\nu = 1$ and the x's are for economies generated with $\nu = 3$. Here we see that in general, increasing α dramatically decreases price/wealth variation (note that the price ratio is plotted on a log scale). This justifies the intuition that as α is increased, more "economic equality" is introduced in the form of less preferential bias in the formation of new edges. The approximately linear relationship suggests that the decrease in variation is exponential in α. Furthermore, the data for $\nu = 1$ shows much larger variation, suggesting that a larger value of ν also has the effect of equalizing buyer opportunities and therefore prices.

Acknowledgments

We are grateful to Tejas Iyer and Vijay Vazirani of Georgia Tech for generously providing their software implementing the Devanur et al. [2002] algorithm. We also give many thanks to Nick Montfort for helpful discussions. Robin Pemantle acknowledges the support of NSF grant DMS-0103635.

References

Kenneth J. Arrow and Gerard Debreu. Existence of an equilibrium for a competitive economy. *Econometrica*, 22(3):265–290, July 1954.

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

B. Bollobas. *Random Graphs*. Cambridge University Press, 2001.

Nikhil R. Devanur, Christos H. Papadimitriou, Amin Saberi, and Vijay V. Vazirani. Market equilibrium via a primal-dual-type algorithm. In *Proceedings of the 43rd Annual IEEE Symposium on Foundations of Computer Science*, 2002.

Irving Fisher. PhD thesis, Yale University, 1891.

D. Gale. *Theory of Linear Economic Models*. McGraw Hill, N.Y., 1960.

Matthew Jackson. A survey of models of network formation: Stability and efficiency. In *Group Formation in Economics: Networks, Clubs and Coalitions*. Cambridge University Press, 2003. Forthcoming.

S. Kakade, M. Kearns, and L. Ortiz. Graphical economics. 2004. Preprint.

M. Mitzenmacher. A brief history of generative models for power law and lognormal distributions. *Internet Mathematics*, 1, 2003.