

Graphical Economics

Sham M. Kakade, Michael Kearns, and Luis E. Ortiz

Department of Computer and Information Science
University of Pennsylvania, Philadelphia, PA 19104

Abstract: We introduce a graph-theoretic generalization of classical Arrow-Debreu economics, in which an undirected graph specifies which consumers or economies are permitted to engage in direct trade, and the graph topology may give rise to local variations in the prices of commodities. Our main technical contributions are: (1) a general existence theorem for *graphical equilibria*, which require *local* markets to clear; (2) an improved algorithm for computing approximate equilibria in standard (non-graphical) economies, which generalizes the algorithm of Deng et al. [2002] to non-linear utility functions; (3) an algorithm for computing equilibria in the graphical setting, which runs in time polynomial in the number of consumers in the special but important case in which the graph is a tree (again permitting non-linear utility functions). We also highlight many interesting learning problems that arise in our model, and relate them to learning in standard game theory and economics, graphical games, and graphical models for probabilistic inference.

1 Introduction

Models for the exchange of goods and their prices in a large economy have a long and storied history within mathematical economics, dating back more than a century to the work of Walras [1874] and Fisher [1891], and continuing through the model of Wald [1936] (see also Brainard and Scarf [2000]). A pinnacle of this line of work came in 1954, when Arrow and Debreu provided extremely general conditions for the existence of an equilibrium in such models (in which markets clear, *i.e.* supply balances demand, and all individual consumers and firms optimize their utility subject to budget constraints). Like Nash’s roughly contemporary proof of the existence of equilibria for normal-form games (Nash [1951]), Arrow and Debreu’s result placed a rich class of economic models on solid mathematical ground.

These important results established the *existence* of various notions of equilibria. The *computation* of game-theoretic and economic equilibria has been a more slippery affair. Indeed, despite decades of effort, the computational complexity of computing a Nash equilibrium for a general-sum normal-form game remains unknown, with the best known algorithms requiring exponential time in the worst case. Even less is known regarding the computation of Arrow-Debreu equilibria. Only quite recently, a polynomial-time algorithm was discovered for the special but challenging case of linear utility functions (Devanur et al. [2002], Jain et al. [2003], Devanur and Vazirani [2003]). Still less is known about the *learning* of economic equilibria in a distributed, natural fashion.

One promising direction for making computational progress is to introduce alternative ways of *representing* these problems, with the hope that wide classes of “natural” problems may permit special-purpose solutions. By developing new representations that permit the expression of common types of structure in games and economies, it may be possible to design algorithms that exploit this structure to yield computational as well

as modeling benefits. Researchers in machine learning and artificial intelligence have proven especially adept at devising models that balance representational power with computational tractability and learnability, so it has been natural to turn to these literatures for inspiration in strategic and economic models.

Among the most natural and common kinds of structure that arise in game-theoretic and economic settings are constraints and asymmetries in the *interactions* between the parties. By this we mean, for example, that in a large-population game, not all players may directly influence the payoffs of all others. The recently introduced formalism of *graphical games* captures this notion, representing a game by an undirected graph and a corresponding set of local game matrices (Kearns et al. [2001]). In Section 2 we briefly review the history of graphical games and similar models, and their connections with other topics in machine learning and probabilistic inference.

In the same spirit, in this paper we introduce a new model called *graphical economics* and show that it provides representational and algorithmic benefits for Arrow-Debreu economics. Each vertex i in an undirected graph represents an individual party in a large economic system. The presence of an edge between i and j means that free trade is allowed between the two parties, while the absence of this edge means there is an embargo or other restriction on direct trade. The graph could thus represent a network of individual business people, with the edges indicating who knows whom; or the global economy, with the edges representing nation pairs with trade agreements; and many other settings. Since not all parties may directly engage in trade, the graphical economics model permits (and realizes) the emergence of *local* prices — that is, *the price of the same good may vary* across the economy. Indeed, one of our motivations in introducing the model is to capture the fact that price differences for identical goods can arise due to the network structure of economic interaction.

We emphasize that the mere introduction of a network or graph structure into economic models is in itself not a new idea; while a detailed history of such models is beyond our scope, Jackson [2003] provides an excellent survey. However, to our knowledge, the great majority of these models are designed to model specific economic settings. Our model has deliberately incorporated a network model into the general Arrow-Debreu framework. Our motivation is to capture and understand network interactions in what is the most well-studied of mathematical economic models.

The graphical economics model suggests a *local* notion of clearance, directly derived from that of the Arrow-Debreu model. Rather than asking that the entire (global) market clear in each good, we can ask for the stronger “provincial” conditions that the *local* market for each good must clear. For instance, the United States is less concerned that the worldwide production of beef balances worldwide demand than it is that the production of *American* beef balances *worldwide* demand for American beef. If this latter condition holds, the American beef industry is doing a good job at matching the global demand for their product, even if other countries suffer excess supply or demand.

The primary contributions of this paper are:

- The introduction of the graphical economics model (which lies within the Arrow-Debreu framework) for capturing structured interaction between individuals, organizations or nations.
- A proof that under very general conditions (essentially analogous to Arrow and Debreu’s original conditions), graphical equilibria always exist. This proof requires a non-trivial modification to that of Arrow and Debreu.

- An algorithm for computing approximate standard market equilibria in the non-graphical setting that runs in time polynomial in the number of players (fixing the number of goods) for a rather general class of non-linear utility functions. This result generalizes the algorithm of Deng et al. [2002] for linear utility functions.
- An algorithm, called **ADProp** (for *Arrow-Debreu Propagation*) for computing approximate graphical equilibria. This algorithm is a message-passing algorithm working directly on the graph, in which neighboring consumers or economies exchange information about trade imbalances between them under potential equilibria prices. In the case that the graph is a tree, the running time of the algorithm is exponential in the graph degree and number of goods k , but only polynomial in the number of vertices n (consumers or economies). It thus represents dramatic savings over treating the graphical case with a non-graphical algorithm, which results in a running time exponential in n (as well as in k).
- A discussion of the many challenging learning problems that arise in both the traditional and graphical economic models. This discussion is provided in Section 6.

2 A Brief History of Graphical Games

In this section, we review the short but active history of work in the model known as *graphical games*, and highlight connections to more longstanding topics in machine learning and graphical models.

Graphical games were introduced in Kearns et al. [2001], where a representation consisting of an undirected graph and a set of local payoff matrices was proposed for multi-player games. The interpretation is that the payoff to player i is a function of the actions of only those players in the neighborhood of vertex i in the graph. Exactly as with the graphical models for probabilistic inference that inspired them (such as Bayesian and Markov networks), graphical games provide an exponentially more succinct representation in cases where the number of players is large, but the degree of the interaction graph is relatively small.

A series of papers by several authors established the computational benefits of this model. Kearns et al. [2001] gave a provably efficient (polynomial in the model size) algorithm for computing all approximate Nash equilibria in graphical games with a tree topology; this algorithm can be formally viewed as the analogue of the junction tree algorithm for inference in tree-structured Markov networks. A related algorithm described in Littman et al. [2002] computes a single but exact Nash equilibrium.

In the same way that the junction tree and polytree algorithms for probabilistic inference were generalized to obtain the more heuristic belief propagation algorithm, Ortiz and Kearns [2003] proposed the NashProp algorithm for arbitrary graphical games, proved its convergence, and experimentally demonstrated promising performance on a wide class of graphs. Vickrey and Koller [2002] proposed and experimentally compared a wide range of natural algorithms for computing equilibria in graphical games, and quite recently Blum et al. [2003] developed an interesting new algorithm based on continuation methods.

An intriguing connection between graphical games and Markov networks was established in Kakade et al. [2003], in the context of the generalization of Nash equilibria known as *correlated equilibria*. There it was shown that if G is the underlying graph

of a graphical game, then all the correlated equilibria of the game (up to payoff equivalence) can be represented as a Markov network whose underlying graph is almost identical to G — in particular, only a small number of highly localized connections need to be added. This result establishes a natural and very direct relationship between the *strategic* structure of interaction in a multi-player game, and the *probabilistic* dependency structure of any (correlated) equilibrium. In addition to allowing one to establish non-trivial independencies that must hold at equilibrium, this result is also thought-provoking from a learning perspective, since a series of recent papers has established that correlated equilibrium appears to be the natural convergence notion for a wide class of “rational” learning dynamics. We shall return to this topic when we discuss learning in Section 6.

3 Graphical Economies

The classical Arrow-Debreu (AD in the sequel) economy (without firms) consists of n consumers who trade k commodities of goods amongst themselves in an unrestricted manner. In an AD economy, each unit of commodity $h \in \{1, \dots, k\}$ can be bought by any consumer at prices p_h . We denote the vector of prices to be $\mathbf{p} \in \mathcal{R}_+^k$ (where $\mathcal{R}_+ = \{x \geq 0\}$).

Each consumer i purchases a *consumption plan* $\mathbf{x}^i \in \mathcal{R}_+^k$, where x_h^i is the amount of commodity h that is purchased by i . We assume that each consumer i has an initial *endowment* $\mathbf{e}^i \in \mathcal{R}_+^k$ of the k commodities, where e_h^i is the amount of commodity h initially held by i . These commodities can be sold to other consumers and thus provide consumer i with *wealth* or cash, which can in turn be used to purchase other goods. Hence, if the initial endowment of consumer i is completely sold, then the wealth of consumer i is $\mathbf{p} \cdot \mathbf{e}^i$. A consumption plan \mathbf{x}^i is *budget constrained* if $\mathbf{p} \cdot \mathbf{x}^i \leq \mathbf{p} \cdot \mathbf{e}^i$, which implicitly assumes the endowment is completely sold (which in fact holds at equilibrium).

Every consumer i has a *utility function* $u_i : \mathcal{R}_+^k \rightarrow \mathcal{R}_+$, where $u_i(\mathbf{x}^i)$ describes how much utility consumer i receives from consuming the plan \mathbf{x}^i . The utility function thus expresses the preferences a consumer has for varying bundles of the k goods.

A *graphical economy* with n players and k goods can be formalized as a standard AD economy with nk “traditional” goods, which are indexed by the pairs (i, h) . The good (i, h) is interpreted as “good h sold by consumer i ”. The key restriction is that free trade is not permitted between consumers, so all players may not be able to purchase (i, h) . It turns out that with these trade restrictions, we were not able to invoke the original existence proof used in the standard Arrow-Debreu model, and we had to use some interesting techniques to prove existence.

It is most natural to specify the trade restrictions through an undirected graph, G , over the n consumers¹. The graph G specifies how the consumers are allowed to trade with each other — each consumer may have a limited choice of where to purchase commodities. The interpretation of G is that if (i, j) is an edge in G , then free trade

¹ Throughout the paper we describe the model and results in the setting where the graph constrains exchange between individual consumers, but everything generalizes to the case in which the vertices are themselves complete AD economies, and the graph is viewed as representing trade agreements.

exists between consumers i and j , meaning that i is allowed to buy commodities from j and vice-versa; while the lack of an edge between i and j means that no direct trade is permitted. More precisely, if we use $N(i)$ to denote the neighbor set of i (which by convention includes i itself), then consumer i is free to buy any commodity *only* from any of the consumers in $N(i)$. It will naturally turn out that rational consumers only purchase goods from a neighbor with the best available price.

Associated with each consumer i is a *local price vector* $\mathbf{p}^i \in \mathcal{R}_+^k$, where p_h^i is the price at which commodity h is being sold by i . We denote the set of all local price vectors by $P = \{\mathbf{p}^i : i = 1, \dots, n\}$. Each consumer i purchases an amount of commodities $\mathbf{x}^{ij} \in \mathcal{R}_+^k$, where x_h^{ij} is the amount of commodity h that is purchased from consumer j by consumer i . The trade restrictions imply that $\mathbf{x}^{ij} = 0$ for $j \notin N(i)$. Here, the consumption plan is the set $X^i = \{\mathbf{x}^{ij} : j \in N(i)\}$ and an X^i is *budget constrained* if $\sum_{j \in N(i)} \mathbf{p}^j \cdot \mathbf{x}^{ij} \leq \mathbf{p}^i \cdot \mathbf{e}^i$ which again implicitly assumes the endowment is completely sold (which holds at equilibrium).

In the graphical setting, we assume the utility function only depends on the *total amount* of each commodity consumed, independent of whom it was purchased from. This expresses the fact that the goods are identical across the economy, and consumers seek the best prices available to them. Slightly abusing notation, we define $\mathbf{x}^i = \sum_{j \in N(i)} \mathbf{x}^{ij}$, which is the total vector amount of goods consumed by i under the plan X^i . The utility of consumer i is given by the function $u_i(\mathbf{x}^i)$, which is a function from $\mathcal{R}_+^k \rightarrow \mathcal{R}_+$.

4 Graphical Equilibria

In equilibrium, there are two properties which we desire to hold — consumer rationality and market clearance. We now define these and state conditions under which an equilibrium is guaranteed.

The economic motivation for a consumer in the choice of consumption plans is to maximize utility subject to a budget constraint. We say that a consumer i uses an *optimal plan* at prices P if the plan maximizes utility over the set of all plans which are budget constrained under P . For instance, in the graphical setting, a plan X^i for i is optimal at prices P if the plan X^i maximizes the function u_i over all X^i subject to $\sum_{j \in N(i)} \mathbf{p}^j \cdot \mathbf{x}^{ij} \leq \mathbf{p}^i \cdot \mathbf{e}^i$.

We say the *market clears* if the supply equals the demand. In the standard setting, define the total demand vector as $\mathbf{d} = \sum_i \mathbf{x}^i$ and the total supply vector as $\mathbf{e} = \sum_i \mathbf{e}^i$ and say the market clears if $\mathbf{d} = \mathbf{e}$. In the graphical setting, the concept of clearance is applied to each “commodity h sold by i ”, so we have a *local* notion of clearance, in which all the goods sold by each consumer clear in the neighborhood. Define the local demand vector $\mathbf{d}^i \in \mathcal{R}_+^k$ on consumer i as $\mathbf{d}^i = \sum_{j \in N(i)} \mathbf{x}^{ji}$. The clearance condition is for each i , $\mathbf{d}^i = \mathbf{e}^i$.

A *market or graphical equilibrium* is a set of prices and plans in which all plans are optimal at the current prices and in which the market clears. We note that the notions of traditional AD and graphical equilibria coincide when the graph is fully connected.

As with the original notion of AD equilibria, it is important to establish the general existence of graphical equilibria. Also as with the original notion, in order to prove the existence of equilibria, two natural technical assumptions are required, one on the utility functions and the other on the endowments. We begin with the assumption on utilities.

Assumption I: For all consumers i , the utility function u_i satisfies the following three properties:

- (Continuity) u_i is a continuous function.
- (Monotonicity) u_i is strictly monotonically increasing with each commodity.
- (Quasi-Concavity) If $u_i(x') > u_i(x)$ then $u_i(\alpha x' + (1 - \alpha)x) > u_i(x)$ for all $0 < \alpha < 1$.

The monotonicity assumption is somewhat stronger than the original “non-satiability” assumption made by AD, but is made primarily for expository purposes. Our results can be generalized to the original assumption as well.

The following facts follow from Assumption I and the consumers’ rationality:

1. At equilibrium, the budget constraint inequality for consumer i is saturated, *e.g.*, in a standard AD economy, a consumer using an equilibrium plan x^i spends all the money obtained from the sale of the endowment e^i .
2. In any graphical equilibrium, a consumer only purchases a commodity at the cheapest price among the neighboring consumers. Note that the neighboring consumer with the cheapest price may not be unique.

Assumption II: (Non-Zero Endowments) For each consumer i and good h , $e_h^i > 0$.

The seminal theorem of Arrow and Debreu [1954] states that these assumptions are sufficient to ensure existence of a market equilibrium. However, this theorem does not immediately imply existence of an equilibrium in a graphical economy, due to the restricted nature of trade. Essentially, Assumption II in the AD setting implies that each consumer owns a positive amount of every good in the economy. In the graphical setting, there are effectively nk goods, but each consumer only has an endowment in k of them. To put it another way, consumer i may only obtain income from selling goods at the k local prices p^i , and is *not* able to sell any of its endowment at prices p^j for $j \neq i$.

Nevertheless, Assumptions I and II still turn out to be sufficient to allow us to prove the following graph-theoretic equilibrium existence theorem.

Theorem 1. (*Graphical Equilibria Existence*) For any graphical economy in which Assumptions I and II hold, there exists a graphical equilibrium.

Before proving existence, let us examine these equilibria with some examples.

4.1 Local Price Variation at Graphical Equilibrium

To illustrate the concept of graphical equilibrium and its difference with the traditional AD notion, we now provide an example in which local price differences occur at equilibrium. The economy consists of three consumers, c_1 , c_2 and c_3 , and two goods, g_1 and g_2 . The graph of the economy is the line $c_1 - c_2 - c_3$.

The utility functions for all three consumers are linear. Consumer c_1 has linear utility for g_1 with coefficient 1, and zero utility for g_2 . Consumer c_2 has linear utility for both g_1 and g_2 , with both coefficients 1. Consumer c_3 , has zero utility for g_1 , and linear utility for g_2 with coefficient 1. The endowments (e_1, e_2) for g_1 and g_2 for the consumers are as follows: $(1, 2)$ for c_1 , $(1, 1)$ for c_2 , and $(2, 1)$ for c_3 .

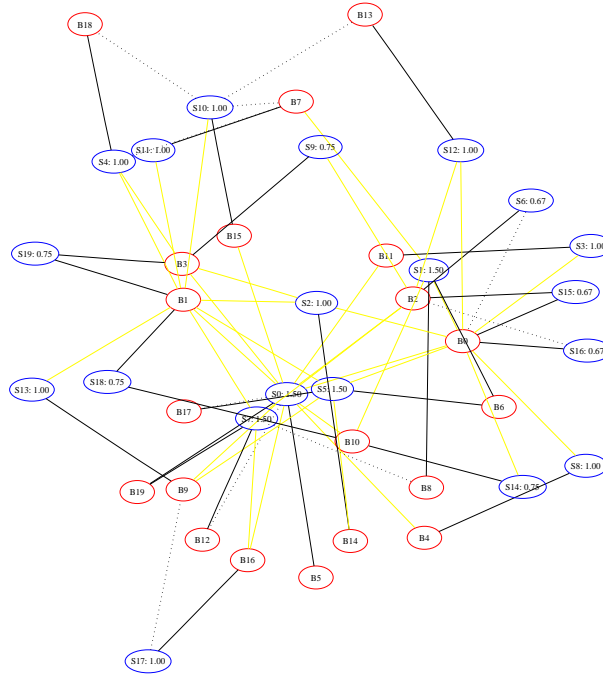


Fig. 1. Price variation and the exchange subgraph at graphical equilibrium in a preferential attachment network. See text for description.

We claim that the following local prices (p_1, p_2) for g_1 and g_2 constitute a graphical equilibrium: prices $(2, 1)$ to purchase from c_1 , $(2, 2)$ to purchase from c_2 , and $(1, 2)$ to purchase from c_3 . It can also be shown that there is *no* graphical equilibrium in which the prices for both goods is the same from all consumers, so price variations are essential for equilibrium. We leave the verification of these claims as an exercise for the interested reader.

Essentially, in this example, c_1 and c_3 would like to exchange goods, but the graphical structure prohibits direct trade. Consumer c_2 , however, is indifferent to the two goods, and thus acts as a kind of arbitrage agent, selling each of c_1 and c_2 their desired good at a high price, while buying their undesired good at a low price.

A more elaborate and interesting equilibrium computation which also contains price variation is shown in Figure 4.1. In this graph, there are 20 buyers and 20 sellers (labeled by ‘B’ or ‘S’ respectively, followed by an index). The bipartite connectivity structure (in which edges are only between buyers and sellers) was generated according to a statistical model known as *preferential attachment* (Barabasi and Albert [1999]), which accounts for the heavy-tailed distribution of degrees often found in real social and economic networks. All buyers have a single unit of currency and utility only for an abstract good, while all sellers have a single unit of this good and utility only for currency. Each seller vertex is labeled with the price they charge at graphical equilibrium. Note that in this example, there is non-trivial price variation, with the most fortunate sellers charging 1.50 for the unit of the good, and the least fortunate 0.67.

The black edges in the figure show the *exchange subgraph* — those pairs of buyers and sellers who actually exchange currency and goods at equilibrium. Note the sparseness of this graph compared to the overall graph. The yellow edges (the most faint in a black and white version) are edges of the original graph that are unused at equilibrium because they represent inferior prices for the buyers, while the dashed edges are edges of the original graph that have competitive prices, but are unused at equilibrium due to the local market clearance conditions.

In a forthcoming paper (Kakade et al. [2004]) we report on a series of large-scale computational experiments of this kind.

4.2 Proof of Graphical Equilibrium Existence

For reasons primarily related to Assumption II, the proof uses the interesting concept of a “quasi-equilibrium”, originally defined by Debreu [1962] in work a decade after his seminal existence result with Arrow. It turns out that much previous work has gone into weakening this assumption in the AD setting. If this assumption is not present, then Debreu [1962] shows that although true equilibria may not exist, “quasi-equilibrium” still exist. In a quasi-equilibrium, consumers with 0 wealth are allowed to be irrational.

Our proof proceeds by establishing the existence of a quasi-equilibria in the graphical setting, and then showing that this in fact implies existence of graphical equilibria. This last step involves a graph-theoretic argument showing that every consumer has positive wealth.

A “graphical quasi-equilibrium” is defined as follows.

Definition 1. A graphical quasi-equilibrium for a graphical economy is a set of globally normalized prices P (i.e. $\sum_{i,h} p_h^i = 1$) and a set of consumption plans $\{X^i\}$, in which the local markets clear and for each consumer i , with wealth $w^i = \mathbf{p}^i \cdot \mathbf{e}^i$, the following condition holds:

- (Rational) If consumer i has positive wealth ($w^i > 0$), then i is rational (utility-maximizing).
- (Quasi-Rational) Else if has no wealth ($w^i = 0$), then the plan X^i is only budget constrained (and does not necessarily maximize utility).

Lemma 1. (Graphical Quasi-Equilibria Existence) In any graphical economy in which Assumption I holds, there exists a graphical quasi-equilibrium.

The proof is straightforward and is provided in a longer version of this paper. Note that if all consumers have positive wealth at a quasi-equilibrium, then all consumers are rational. Hence, to complete the proof of Theorem 1 it suffices to prove that all consumers have positive wealth at a quasi-equilibrium. For this we provide the following lemma, which demonstrates how wealth propagates in the graph.

Lemma 2. If the graph of a graphical economy is connected and if Assumptions I and II hold, then for any quasi-equilibrium set of prices $\{\mathbf{p}^i\}$, it holds that every consumer has non-zero wealth.

Proof. Note that by price normalization, there exists at least one consumer that has one commodity with non-zero price. We now show that if for any consumer i , $\mathbf{p}^i \neq \mathbf{0}$, then this implies that for all $j \in N(i)$, $\mathbf{p}^j \neq \mathbf{0}$. This is sufficient to prove the result, since the graph is assumed to be connected and $e^i > 0$.

Let $\{X^i\}$ and $\{\mathbf{p}^i\}$ be a quasi-equilibrium. Assume that in some i , $\mathbf{p}^i \neq \mathbf{0}$. Since every consumer has positive endowments in each commodity (Assumption II), $\mathbf{p}^i \cdot \mathbf{e}^i > 0$, and so consumer i is rational. By Fact 1, the budget constraint inequality of i must be saturated, so $\sum_{j \in N(i)} \mathbf{p}^j \cdot \mathbf{x}^{ij} = \mathbf{p}^i \cdot \mathbf{e}^i > 0$. Hence, there must exist a commodity h and a $j \in N(i)$ such that $x_h^{ij} > 0$ and $p_h^j \neq 0$, else the money spent would be 0. In other words, there must exist a commodity that is consumed by i from a neighbor at a non-zero price.

The rationality of i implies that consumer j has the cheapest price for the commodity h , otherwise i would buy h from a cheaper neighbor (Fact 2). More formally, $j \in \arg \min_{\ell \in N(i)} p_h^\ell$, which implies for all $\ell \in N(i)$, $p_h^\ell \geq p_h^j > 0$. Thus we have shown that for all $\ell \in N(i)$, $\mathbf{p}^\ell \neq \mathbf{0}$, and since by Assumption II, $e^\ell > 0$, this completes the proof. \square

Without graph connectivity, it is possible that all the consumers in a disconnected graph could have zero wealth at a quasi-equilibrium. Hence, to complete the proof of Theorem 1, we observe that in each connected region we have a separate graphical equilibria.

It turns out that the ‘‘propagation’’ argument in the previous proof, with more careful accounting, actually leads to a quantitative lower bound on consumer wealth in a graphical economy, which we now present. This lower bound is particularly useful when we turn towards computational issues in a moment.

The following definitions are needed:

$$e_+ = \max_{i,h} e_h^i, \quad e_- = \min_{i,h} e_h^i$$

Note that Assumption II implies that $e_- > 0$.

Lemma 3. (*Wealth Propagation*) *In a graphical economy, in which Assumptions I and II hold, with a connected graph of degree $m - 1$, the wealth of any consumer i at equilibrium prices $\{\mathbf{p}^i\}$ is bounded as follows:*

$$\mathbf{p}^i \cdot \mathbf{e}^i \geq \left(\frac{e_-}{e_+ m k} \right)^{\text{diameter}(G)} \frac{e_-}{n} > 0$$

The proof is provided in the long version of this paper.

Interestingly, note that a graph that maximizes free trade (*i.e.* a fully connected graph) maximizes this lower bound on the wealth of a consumer.

5 Algorithms for Computing Economic Equilibria

All of our algorithmic results compute approximate, rather than exact, economic equilibria. We first give the requisite definitions. We use the natural definition originally presented in Deng et al. [2002]. First, two concepts are useful to define — approximate optimality and approximate clearance. A plan is ε -optimal at some price P if the plans

are budget constrained under P and if the utility of the plan is at least $1 - \varepsilon$ times the optimal utility under P . The market ε -clears if, in the standard setting, $(1 - \varepsilon)e \leq \mathbf{d} \leq \mathbf{e}$ and, in the graphical setting, for all i , $(1 - \varepsilon)e^i \leq \mathbf{d}^i \leq e^i$. Now we say a set of plans and prices constitute an ε -equilibrium if the market ε -clears and if the plans are ε -optimal.²

The algorithms we present search for an approximate ADE on a discretized grid. Hence, we need some sort of “smoothness” condition on the utility function in order for the discretized grid to be a good approximation to the true space. More formally,

Assumption III We assume there exists $\gamma \geq 0$ such that for all i and for all \mathbf{x}

$$u_i((1 + \gamma)\mathbf{x}) \leq \exp(\gamma d)u_i(\mathbf{x})$$

for some constant d .

Note that for polynomials with positive weights, the constant d can be taken to be the degree of the polynomial. Essentially, the condition states that if a consumer increases his consumption plan by some multiplicative factor γ , then his utility cannot increase by the exponentially larger, multiplicative factor of $\exp(\gamma d)$. This condition is a natural one to consider, since the “growth rate” constant d is dimensionless (unlike the derivative of the utility function $\partial u_i / \partial \mathbf{x}$, which has units of *utility/goods*).

Naturally, for reasons of computational generality, we make a “black box” representational assumption on the utility functions.

Assumption IV We assume that for all i , the utility function u_i is given as an oracle, which given an input \mathbf{x}^i , outputs $u_i(\mathbf{x}^i)$ in unit time.

For the remainder of the paper, we assume that Assumptions I-IV hold.

5.1 An Improved Algorithm for Computing AD Equilibria

We now present an algorithm for computing AD equilibria for rather general utility functions in the non-graphical setting. The algorithm is a generalization of the algorithm provided by Deng et al. [2002], which computes equilibria for the case in which the utilities are linear functions. While our primary interest in this algorithm is as a subroutine for the graphical algorithm presented in Section 5.3, it is also of independent interest.

The idea of the algorithm is as follows. For each consumer i , a binary valued “best-response” table $M_i(\mathbf{p}, \mathbf{x})$ is computed, where the indices \mathbf{p} and \mathbf{x} are prices and plans. The value of $M_i(\mathbf{p}, \mathbf{x})$ is set to 1 if and only if \mathbf{x} is ε -optimal for consumer i at prices \mathbf{p} . Once these tables are computed, the “price player’s” task is then to find \mathbf{p} and $\{\mathbf{x}^i\}$ such that $(1 - \varepsilon)e \leq \mathbf{d} \leq \mathbf{e}$ and for all i , $M_i(\mathbf{p}, \mathbf{x}^i) = 1$.

To keep the tables of M_i of finite size, we only consider prices and plans on a grid. As in Deng et al. [2002] and Papadimitriou and Yannakakis [2000], we consider a

² It turns out that any ε -approximate equilibrium in our setting with monotonically increasing utility functions can be transformed into an approximate equilibrium in which the market *exactly* clears while the plans are still ε -optimal. To see this note that the cost of the unsold goods is equal to the surplus money in the consumers’ budgets. The monotonicity assumption allows us to increase the consumption plans, using the surplus money, to take up the excess supply without decreasing utilities. This transformation is in general not possible if we weaken the monotonicity assumption to a non-satiability assumption.

relative grid of the form:

$$\mathcal{G}_{\text{price}} = \{p_0, (1 + \varepsilon)p_0, (1 + \varepsilon)^2 p_0, \dots, 1\},$$

$$\mathcal{G}_{\text{plan}} = \{x_0, (1 + \varepsilon)x_0, (1 + \varepsilon)^2 x_0, \dots, ne_+\}$$

where the maximal grid price is 1 and maximal grid plan is ne_+ (since there is at most an amount ne_+ of any good in the market). The intuitive reason for the use of a relative grid is that demand is more sensitive to price perturbations of cheaper priced goods, since consumers have more purchasing power for these goods.

In Section 5.2, we sketch the necessary approximation scheme, which shows how to set p_0 and x_0 such that an ε -equilibrium on this grid exists. The natural method to set p_0 is to use a lower bound on the equilibrium prices. Unfortunately, under rather general conditions, only the trivial lower bound of 0 is possible. However, we can set p_0 and x_0 based on a non-trivial *wealth* bound.

Now let us sketch how we use the tables to compute an ε -equilibrium. Essentially, the task now lies in checking that the demand vector \mathbf{d} is close to \mathbf{e} for a set of plans and prices which are true for the M_i . As in Deng et al. [2002], a form of dynamic programming suffices. Consider a binary, “partial sum of demand” table $S_i(\mathbf{p}, \mathbf{x})$ defined as follows: $S_i(\mathbf{p}, \mathbf{d}) = 1$ if and only if there exists x^1, \dots, x^i such that $\mathbf{d} = x^1 + x^2 + \dots + x^i$ and $M_1(\mathbf{p}, x^1) = 1, \dots, M_i(\mathbf{p}, x^i) = 1$. These tables can be computed recursively as follows: if $S_{i-1}(\mathbf{p}, \mathbf{d}) = 1$ and if $M_i(\mathbf{p}, x) = 1$, then we set $S_i(\mathbf{p}, \mathbf{d} + x) = 1$. Further, we keep track of a “witness” x^1, \dots, x^i which proves that the table entry is 1. The approximation lemmas in Section 5.2 show how to keep this table of finite “small” size (see also long version of the paper).

Once we have S_n , we just search for some index \mathbf{p} and \mathbf{d} such that $S_n(\mathbf{p}, \mathbf{d}) = 1$ and $\mathbf{d} \approx \mathbf{e}$. This \mathbf{p} and the corresponding witness plans then constitute an equilibrium. The time complexity of this algorithm is polynomial in the tables sizes, which we shall see is of polynomial size for a fixed k . This gives rise to the following theorem.

Theorem 2. *For fixed k , there exists an algorithm which takes as input an AD economy and outputs an ε -equilibrium in time polynomial in $n, 1/\varepsilon, \log(e_+/e_-)$, and d .*

The approximation details and proof are provided in the long version of this paper.

5.2 Approximate Equilibria on a Relative Grid

We now describe a relative discretization scheme for prices and consumption plans that is used by the algorithm just described for computing equilibria in classical (non-graphical) AD economies. This scheme can be generalized for the graphical setting, but is easier to understand in the standard setting.

Without loss of generality, throughout this section we assume the prices in a market are globally normalized, *i.e.* $\sum_h p_h = 1$.

A price and consumption plan can be mapped onto the relative grid in the obvious way. Define $grid(\mathbf{p}) \in \mathcal{R}_+^k$ to be the closest price to \mathbf{p} such that each component of $grid(\mathbf{p})$ is on the price grid. Hence,

$$\frac{1}{1 + \varepsilon} \mathbf{p} \leq grid(\mathbf{p}) \leq \max\{(1 + \varepsilon)\mathbf{p}, p_0 \mathbf{1}\}$$

where the max is taken component-wise and $\mathbf{1}$ is a k -length vector of all ones. Note that the value of p_0 is a threshold where all prices below p_0 get set to this threshold price. Similarly, for any consumption plan \mathbf{x}^i , let $grid(\mathbf{x}^i)$ be the closest plan to \mathbf{x}^i such that $grid(\mathbf{x}^i)$ is componentwise on \mathcal{G}_{plan} .

In order for such a discretization scheme to work, we require two properties. First, the grid should certainly *contain* an approximate equilibrium of the desired accuracy. We shall refer to this property as *Approximate Completeness* (of the grid). Second, and more subtly, it should also be the case that maximizing consumer utility, while *constrained* to the grid, results in utilities close to those achieved by the *unconstrained* maximization — otherwise, our grid-restricted search for equilibria might result in highly suboptimal consumer plans. We shall refer to this property as *Approximate Soundness* (of the grid). It turns out that Approximate Soundness only holds if prices ensure a minimum level of wealth for each consumer, but conveniently we shall always be in such a situation due to Lemma 3.

The next two lemmas establish Approximate Completeness and Soundness for the grid. The Approximate Completeness Lemma also states how to set p_0 and x_0 . It is straightforward to show that if we have a lower bound on the price at equilibrium, then p_0 can be set to this lower bound. Unfortunately, it turns out that under our rather general conditions we cannot provide a lower bound. Instead, as the lemmas show, it is sufficient to use a lower bound w_0 on the wealth of any consumer at equilibrium, and set p_0 and x_0 based on this wealth. Note that in the traditional AD setting e_- is a bound on the wealth, since the prices are normalized.

Lemma 4. (*Approximate Completeness*) *Let the grids \mathcal{G}_{price} and \mathcal{G}_{plan} be defined using*

$$p_0 = \frac{\varepsilon}{nke_+}w_0, \quad x_0 = \frac{\varepsilon}{(1+13\varepsilon)nk}w_0$$

where w_0 is a lower bound on equilibrium wealth of all consumers and let $\{\mathbf{x}^{*i}\}$ and $\{\mathbf{p}^{*i}\}$ be equilibrium prices and plans. Then the plans $\{\mathbf{x}^i = grid\left(\frac{1}{1+13\varepsilon}\mathbf{x}^{*i}\right)\}$ are $19d\varepsilon$ approximately optimal for the price $\mathbf{p} = grid(\mathbf{p}^*)$ and the market 14ε -approximately clears. Furthermore, a useful property of this approximate equilibrium is that every consumer has wealth greater than $\frac{w_0}{1+\varepsilon}$.

There are a number of important subtleties to be addressed in the proof, which we formally present in the longer version. For instance, note that the closest point on the grid to some true equilibria may not even be budget constrained.

Lemma 5. (*Approximate Soundness*) *Let the grid be defined as in Theorem 4 and let \mathbf{p} be on the grid such that every consumer has wealth above $\frac{w_0}{1+\varepsilon}$. If the plans $\{\mathbf{x}^i\}$ β -approximately maximize utility over the budget constrained plans which are componentwise on the grid, i.e. if for all budget constrained \mathbf{x}^{*i} which lie on the plan grid,*

$$u_i(\mathbf{x}^i) \geq (1 - \beta)u_i(\mathbf{x}^{*i}).$$

then

$$u_i(\mathbf{x}^i) \geq (1 - (\beta + 4\varepsilon d))u_i^*$$

where u_i^* is the optimal utility under \mathbf{p} .

5.3 Arrow-Debreu Propagation for Graphical Equilibria

We now turn to the problem of computing equilibria in graphical economies. We present the **ADProp** algorithm, which is a dynamic programming, message-passing, algorithm for computing approximate graphical equilibria when the graph has a tree structure. Recall that in a graphical economy there are effectively nk goods, so we cannot keep the number of goods fixed as we scale the number of consumers. Hence, the algorithm described in the previous section cannot be directly applied if we wish to scale polynomially with the number of consumers.

As we will see from the description of **ADProp** below, an appealing conceptual property of the algorithm is how it achieves the computation of *global* economic equilibria in a distributed manner through the *local* exchange of economic trade and price information between just the neighbors in the graph.

We orient the graph such that “downstream” from a vertex lies the root and “upstream” lies the leaves. For any consumer j that is not the root there exists a unique downstream consumer, say ℓ . Let $UP(j)$ be the set of neighbors of j which are not downstream, *i.e.* $UP(j)$ is the set $N(j) - \{\ell\}$ so it includes j itself.

We now define a binary valued table $T_{\ell j}$, which can be viewed as the message that consumer $j \in UP(\ell)$ sends downstream to ℓ . The table $T_{\ell j}(\mathbf{p}^\ell, \mathbf{x}^{\ell j}, \mathbf{p}^j, \mathbf{x}^{j\ell})$ is indexed by the prices for ℓ and j and the consumption that flows along the edge between ℓ and j — from ℓ to j , the consumption is $\mathbf{x}^{\ell j}$, and from j to ℓ , the consumption is $\mathbf{x}^{j\ell}$. The table entry $T_{\ell j}(\mathbf{p}^\ell, \mathbf{x}^{\ell j}, \mathbf{p}^j, \mathbf{x}^{j\ell})$ evaluates to 1 if and only if there exists a *conditional* ε -equilibria upstream from j (inclusive) in which the respective prices and plans are fixed to $\mathbf{p}^\ell, \mathbf{x}^{\ell j}, \mathbf{p}^j, \mathbf{x}^{j\ell}$. For the special case where $j = \ell$, the table entry $T_{jj}(\mathbf{p}^\ell, \mathbf{x}^{\ell j}, \mathbf{p}^j, \mathbf{x}^{j\ell})$ is set to 1 if and only if $\mathbf{p}^\ell = \mathbf{p}^j$ and $\mathbf{x}^{\ell j} = \mathbf{x}^{j\ell}$ (note that \mathbf{x}^{jj} is effectively the amount of the goods that j desires not to sell).

The tables provide all the information needed to apply dynamic programming in the obvious way. In its *downstream pass*, **ADProp** computes the table $T_{\ell j}$ recursively, in the typical dynamic programming fashion. If j is an internal node in the tree, when j has received the appropriate tables from all $i \in UP(j)$, we must set $T_{\ell j}(\mathbf{p}^\ell, \mathbf{x}^{\ell j}, \mathbf{p}^j, \mathbf{x}^{j\ell}) = 1$, if: 1) a conditional upstream equilibrium exists, which we can compute from the tables passed to j , 2) the plan X^j , consistent with the upstream equilibrium, is ε -optimal for the neighborhood prices, and 3) the market ε -clear at j . Naturally, a special but similar operation occurs at the leaves and the root of the tree.

Once **ADProp** computes the message at the root consumer, it performs an *upstream pass* to obtain a single graphical equilibrium, again, in the typical dynamic programming fashion. At every node, starting with the root, **ADProp** selects price and allocation assignments consistent with the tables at the node and passes those assignments up to their upstream neighbors, until it reaches the leaves of the tree.

As presented in Section 5.2, we can control the approximation error by using appropriate sized grids. This leads to our main theorem for computing graphical equilibrium.

Theorem 3. (*ADProp*) For fixed k and graph degree, **ADProp** takes as input a tree graphical economy in which Assumptions I-IV hold and outputs an ε -equilibrium in time polynomial in $n, 1/\varepsilon, \log(e_+/e_-)$, and d .

Heuristic generalizations of **ADProp** are possible to handle more complex (loopy) graph structures (*a la* **NashProp** [Ortiz and Kearns, 2003]).

6 Learning in Graphical Games and Economics

Although the work described here has focused primarily on the graphical economics representation, and algorithms for equilibrium computation, the general area of graphical models for economic and strategic settings is rich with challenging learning problems and issues. We conclude by mentioning just a few of these.

Rational Learning in Graphical Games. What happens if each player in a repeated graphical game plays according to some “rational” dynamics (like fictitious play, best response, or other variants), but using only *local* observations (the actions of neighbors)? In cases where convergence occurs, how does the graph structure influence the equilibrium chosen? Are there particular topological properties that favor certain players in the network?

No-Regret Learning in Graphical Games. It has recently been established that if all players in a repeated graphical game play a local *no internal regret* algorithm, the population empirical play will converge to the set of *correlated equilibria*. It was also noted in the introduction that all such equilibrium can be represented up to payoff equivalence on a related Markov network; under what conditions will no-regret learning dynamics actually settle on one of these *succinct* equilibria? In preliminary experiments using the algorithms of Foster and Vohra [1999] as well as those of Hart and Mas-Colell [200] and Hart and Mas-Colell [2001], one does not observe convergence to the set of payoff-equivalent Markov network correlated equilibria.

Learning in Traditional AD Economies. Even in the non-graphical Arrow-Debreu setting, little is known about reasonable distributed learning procedures. Aside from a strong (impossibility) result by Saari and Simon [1978] suggesting that general convergence results may not be possible, there is considerable open territory here. Conceptual challenges include the manner in which the “price player” should be modeled in the learning process.

Learning in Graphical Economics. Finally, problems of learning in the graphical economics model are entirely open, including the analogues to all of the questions above. Generally speaking, one would like to formulate reasonable procedures for *local* learning (adjustment of seller prices and buyer purchasing decisions), and examine how these procedures are influenced by network structure.

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.
- Ben Blum, Christian R. Shelton, and Daphne Koller. A continuation method for Nash equilibria in structured games. In *Proceedings of the Eighteenth International Joint Conference on Artificial Intelligence*, 2003.
- William C. Brainard and Herbert E. Scarf. How to compute equilibrium prices in 1891. *Cowles Foundation Discussion Paper 1272*, August 2000.
- Gerard Debreu. New concepts and techniques for equilibrium analysis. *International Economic Review*, 3(3):257–273, September 1962.
- Xiaotie Deng, Christos Papadimitriou, and Shmuel Safra. On the complexity of equilibria. In *Proceedings of the Thiry-fourth Annual ACM Symposium on Theory of Computing*, pages 67–71. ACM Press, 2002.

- 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.
- Nikhil R. Devanur and Vijay V. Vazirani. An improved approximation scheme for computing Arrow-Debreu prices for the linear case. In *Proceedings of the 23rd Conference on Foundations of Software Technology and Theoretical Computer Science*, 2003. To appear.
- Irving Fisher. PhD thesis, Yale University, 1891.
- D. Foster and R. Vohra. Regret in the on-line decision problem. *Games and Economic Behavior*, pages 7 – 36, 1999.
- Sergiu Hart and Andreu Mas-Colell. A simple adaptive procedure leading to correlated equilibrium. *Econometrica*, 68(5):1127 – 1150, 2000.
- Sergiu Hart and Andreu Mas-Colell. A reinforcement procedure leading to correlated equilibrium. In Gerard Debreu, Wilhelm Neuefeind, and Walter Trockel, editors, *Economic Essays*, pages 181 – 200. Springer, 2001.
- 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.
- Kamal Jain, Mohammad Mahdian, and Amin Saberi. Approximating market equilibria. In *Proceedings of the 6th. International Workshop on Approximation Algorithms for Combinatorial Optimization Problems*, 2003.
- S. Kakade, M. Kearns, J. Langford, and L. Ortiz. Correlated equilibria in graphical games. In *Proceedings of the 4th ACM Conference on Electronic Commerce*, pages 42–47, 2003.
- S. Kakade, M. Kearns, L. Ortiz, R. Pemantle, and S. Suri. The economics of social networks. 2004. Preprint.
- M. Kearns, M. Littman, and S. Singh. Graphical models for game theory. In *Proceedings of the Conference on Uncertainty in Artificial Intelligence*, pages 253–260, 2001.
- M. Littman, M. Kearns, and S. Singh. An efficient exact algorithm for singly connected graphical games. In *Neural Information Processing Systems*, 2002.
- J. F. Nash. Non-cooperative games. *Annals of Mathematics*, 54:286–295, 1951.
- L. Ortiz and M. Kearns. Nash propagation for loopy graphical games. In *Neural Information Processing Systems*, 2003.
- Christos H. Papadimitriou and Mihalis Yannakakis. On the approximability of trade-offs and optimal access of web sources. In *Proceedings of the 41st Annual IEEE Symposium on Foundations of Computer Science*, 2000.
- D. G. Saari and C. P. Simon. Effective price mechanisms. *Econometrica*, 46, 1978.
- D. Vickrey and D. Koller. Multi-agent algorithms for solving graphical games. In *Proceedings of the National Conference on Artificial Intelligence*, 2002.
- Abraham Wald. Über einige Gleichungssysteme der mathematischen Ökonomie (On some systems of equations of mathematical economics). *Zeitschrift für Nationalökonomie*, 7(5):637–670, 1936. English translation by Otto Eckstein in *Econometrica*, Vol. 19, No. 4 (Oct., 1951), 368–403.
- Léon Walras. *Éléments d'économie politique pure; ou, Théorie de la richesse sociale (Elements of Pure Economics, or the Theory of social wealth)*. Lausanne, Paris, 1874. (1899, 4th ed.; 1926, rev ed., 1954, Engl. transl.).