

Graphical Economics with Resale

Gabriel P. Andrade

Department of Computer Science
University of Colorado Boulder
gabriel.andrade@colorado.edu

Sharadha Srinivasan

Department of Computer Science
University of Colorado Boulder
sharadha.srinivasan@colorado.edu

Rafael Frongillo

Department of Computer Science
University of Colorado Boulder
raf@colorado.edu

Elliot Gorokhovskiy

Department of Mathematics
Arizona State University
elliott.gorokhovskiy@gmail.com

ABSTRACT

Kakade, Kearns, and Ortiz (KKO) introduce a graph-theoretic generalization of the classic Arrow–Debreu (AD) exchange economy. Despite its appeal as a networked version of AD, we argue that the KKO model is too local, in the sense that goods cannot travel more than one hop through the network. We introduce an alternative model in which agents may purchase goods on credit in order to resell them. In contrast to KKO, our model allows for long-range trade, and yields equilibria in more settings than KKO, including sparse endowments. Our model smoothly interpolates between the KKO and AD equilibrium concepts: we recover KKO when the resale capacity is zero, and recover AD when it is sufficiently large. We give general equilibrium existence results, and an intuitive auction-based algorithm to compute equilibria when agent utilities satisfy the weak gross-substitutes property.

ACM Reference Format:

Gabriel P. Andrade, Rafael Frongillo, Sharadha Srinivasan, and Elliot Gorokhovskiy. 2021. Graphical Economics with Resale. In *Appears at the 3rd Games, Agents, and Incentives Workshop (GAIW 2021). Held as part of the Workshops at the 20th International Conference on Autonomous Agents and Multiagent Systems., London, UK, May 2021, IFAAMAS*, 18 pages.

1 INTRODUCTION

Market exchange equilibria have an illustrious history dating back to Walras in 1874 [33], culminating in a general proof of existence for Arrow–Debreu (AD) equilibria [1]. A key assumption of the Walrasian framework—that markets are *centralized*—has been under scrutiny the past 20 years, spurring a rapidly growing literature focused on the economics of networks (cf. [5], § 6.4). In these decentralized models, agents are typically nodes in a graph and trade is restricted to edges. A particularly inspiring example, and the prime motivation for the present work, is the graphical exchange economy introduced by Kakade, Kearns, and Ortiz (KKO) [22]. The KKO model allows for *local* markets, where intuitively each agent can set its own prices, and can purchase goods from neighboring agents at their prices.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

Appears at the 3rd Games, Agents, and Incentives Workshop (GAIW 2021). Held as part of the Workshops at the 20th International Conference on Autonomous Agents and Multiagent Systems., Aziz, Ceppi, Dickerson, Hosseini, Lev, Mattei, McElfresh, Zick (chairs), May 2021, London, UK. © 2021 Copyright held by the owner/author(s).

While the KKO model does capture many intuitive phenomena arising from economies with graphical structure, the model falls short in several simple scenarios. In particular, KKO is perhaps “too local”, as the model prohibits goods to move more than one hop in the network. As such, there are simple networks where one would intuitively expect a chain of trade, yet KKO offers either no equilibrium or one which fragments the network. See §3 for two such examples.

In this paper, we introduce an extension of the KKO model which allows agents to purchase goods to resell, thereby facilitating long-range trade throughout the network. As we show, our model provides an intuitive resolution to scenarios described above involving a chain of trade, giving the intermediate nodes a share of the overall surplus. Our model also smoothly interpolates between the KKO and AD equilibrium concepts, by throttling the amount that agents are allowed to resell: when resale is eliminated, goods can only move one hop, and we recover KKO, and conversely when resale is sufficiently large, all local prices must coincide, and we recover AD.

In addition to the intuitive appeal of our model, we give standard existence results, showing that equilibria exist in essentially any setting where AD equilibria exist (§4). We also give an auction based algorithm to compute approximate equilibria, by modifying previous combinatorial auction-based algorithms to account for resale (§5). Despite the similarity to the algorithms in [18, 24], we require entirely new techniques to bound the runtime. Our algorithmic results are perhaps surprising, as our model is closely related to AD economies with production, which in general have much bleaker computational results [14, 16, 17]. We conclude in §6 with extensions, discussion, and future directions.

1.1 Related Work

Understanding the behaviour of markets (economies) is at the heart of many research streams in theoretical economics. Work pioneered by the 19th century economist Léon Walras [33] has been at the epicenter since pioneering work by Arrow and Debreu [1], which establishes the existence of equilibria in a very general model of the economy that is now known as the Arrow–Debreu (AD) market model. In the theoretical computer science community, computing equilibria in the AD model has garnered much attention. In the purely exchange setting of the AD model multiple algorithmic innovations have been made to compute equilibria, e.g. [7, 10, 11, 13, 18, 22, 32]. The work introduced in this paper falls outside the scope of these algorithms since resale in graphical economies behaves like

a special version of production in the AD model. For AD markets with production, Jain and Varadarajan [21] and Kapoor et al. [24] give polynomial-time algorithms for finding competitive equilibria for certain types of production and utility functions. However, resale in graphical economies falls outside the scope of production applicable to the convex program in [21], and is more complicated than the production in [24]. In fact, certain forms of resale reduce to a special case of production where even approximating competitive equilibria is known to be FIXP-hard in general [15, 16].

We introduce a variant of the auction based algorithms by Garg and Kapoor [18] and Kapoor et al. [24] that is able to approximate competitive equilibria in graphical economies with resale, though our proofs differ substantially. By considering specific graph structures, the algorithm being introduced extends both [18] and [24] to a more general class of utility functions and, in the case of [18], to a more general class of production functions. These auction-based algorithms have a storied history in general resource allocation and network optimization [2, 3, 8, 25, 34], thus generalizing these algorithms to a network market model with long range trade provides insights reaching farther than classic economics. Finally, a form of resale discussed in this paper, called credit bound resale, can be thought of as a natural extension of the spending constraint utilities introduced in [10, 32] to a production-like setting.

Along with the aforementioned connections to traditional economic theory and algorithmic economics, the work presented here has deep ties to the study of economic networks more broadly. In §6.4 we discuss these connections in detail.

2 SETTING AND BACKGROUND

We consider economies consisting of a set $[\ell] := \{1, \dots, \ell\}$ of divisible goods and a set $[m] := \{1, \dots, m\}$ of agents embedded as nodes in some graph $G = ([m], E)$, whose edges E describe who may trade with whom. For ease of exposition we will assume that G is undirected throughout this paper; all results can be easily extended to directed graphs. In addition, to simplify notation, in lieu of E we will use the reflexive symmetric binary relation \simeq on $[m]$, so that for any two agents $i, j \in [m]$, the presence of an edge between them is denoted by $i \simeq j$. When it is important to specify, we will also use the notation $i \sim j$ to mean ($i \simeq j$ & $i \neq j$).

In the economy, each agent $i \in [m]$ has an endowment of goods $\mathbf{e}^i \in \mathbb{R}_+^\ell$ and a utility function $u_i : \mathbb{R}_+^\ell \rightarrow \mathbb{R}_+$. The endowment vector \mathbf{e}^i describes the bundle of goods that agent i enters the market with; agent i is endowed an amount e_k^i of good $k \in [\ell]$. The utility function u_i encodes agent i 's preferences over bundles of goods. As introduced by Kakade et al. [22], a graphical economy can be formally defined as follows.

Definition 2.1. A *graphical economy* is an undirected graph G over agents $[m]$ with neighbor relation \simeq , utilities $\{u_i : \mathbb{R}_+^\ell \rightarrow \mathbb{R}_+\}_{i \in [m]}$, and endowments $\{\mathbf{e}^i \in \mathbb{R}_+^\ell\}_{i \in [m]}$, where ℓ is an integer denoting the number of goods being traded.

To discuss equilibria in a graphical economy, Kakade, et al. [22] introduce local price vectors $\mathbf{p}^i \in \mathbb{R}_+^\ell$ for each agent $i \in [m]$, so that p_k^i is the price at which agent i sells one unit of good $k \in [\ell]$.

The consumption plans¹ are given by vectors $\mathbf{x}^{ij} \in \mathbb{R}_+^\ell$ describing the bundle of goods agent $i \in [m]$ purchases from agent $j \in [m]$ for consumption, so that the set $\mathbf{x}^i = \{\mathbf{x}^{ij} : j \in [m]\}$ thus describes all of the consumption for agent i . To enforce the condition that trade must traverse edges, we set $\mathbf{x}^{ij} = \mathbf{0}$ for $j \neq i$.

Throughout the paper we will use $\mathbf{e} = \{\mathbf{e}^i : i \in [m]\}$ to denote the full set of endowments in the economy, $\mathbf{p} = \{\mathbf{p}^i : i \in [m]\}$ to denote the full set of prices in the economy, $\mathbf{x} = \{\mathbf{x}^i : i \in [m]\}$ to denote the full set of consumption plans in the economy, etc. In addition, we will treat each of these sets as vectors, matrices, and tensors that are indexed in a way that is consistent with their definitions above; for example, agent i buys an amount x_k^{ij} of good k from agent j for consumption.

2.1 The Arrow–Debreu (AD) Exchange Model

The Arrow–Debreu (AD) exchange economy [1], also known as the Walrasian model, is extremely well studied due to its central role in general equilibrium theory. The graphical economies studied here are generalizations of AD which retain AD as a special case [22]. In the language introduced above, consider a graphical economy with m agents and ℓ goods that is embedded in a complete graph, i.e. a graphical economy where every pair of agents is able to trade with one another.

Definition 2.2. An *AD equilibrium* is a pair (\mathbf{p}, \mathbf{x}) of a set of price vectors \mathbf{p} and set of consumption plans \mathbf{x} if the underlying graph is complete, we have $\mathbf{p}^i = \mathbf{p}^j$ for all $i, j \in [m]$, and the following conditions are satisfied:

1. *Market Clearing.*

$$\sum_{i, j \in [m]} \mathbf{x}^{ij} = \sum_{i \in [m]} \mathbf{e}^i. \quad (1)$$

2. *Individual rationality (IR).* For all agents $i \in [m]$, setting $\hat{\mathbf{x}}^i = \mathbf{x}^i$ maximizes their utility $u_i(\sum_{j \simeq i} \hat{\mathbf{x}}^{ij})$ over all $\hat{\mathbf{x}}^i \in \mathbb{R}_+^\ell$ satisfying

$$\sum_{j \simeq i} \mathbf{p}^j \cdot \hat{\mathbf{x}}^{ij} \leq \mathbf{p}^i \cdot \mathbf{e}^i. \quad (2)$$

Notice that since the underlying trade graph is complete, the condition that $\mathbf{p}^i = \mathbf{p}^j$ for all $i, j \in [m]$ is without loss of generality: by allowing trade between every pair of agents, at equilibrium the local aspects of a graphical economy essentially become obsolete. To see this, note that agents can consume goods from any other agent, and thus individual rationality dictates that agents should only consume utility maximizing goods at the cheapest price in the economy. To enforce market clearing the price on a good must therefore be identical throughout the economy, meaning that prices can simply be described by a single vector. Similarly, since local prices are all the same, all that matters about the consumption plan \mathbf{x} are the sums $\mathbf{d}^i(\mathbf{x}) = \sum_{j \simeq i} \mathbf{x}^{ji}$, often called the demand vectors, specifying what each agent consumes but not from whom. After

¹In the traditional AD economy setting, an agent's consumption plan is often referred to as their "allocation". However, as defined in §3, we will consider agents who purchase goods for both consumption and reselling. For this reason we opt for the term consumption plan, as it more naturally differentiates between the goods purchased for consumption vs. reselling.

these two simplifications, we arrive at the usual definition of an AD equilibrium.

2.2 The Kakade, Kearns, Ortiz (KKO) Equilibrium

By moving from the complete graph to general graphs, and imposing a local clearing condition instead of a global one, we arrive at the notion of equilibria for graphical economies introduced by Kakade, Kearns, and Ortiz (KKO) [22]. Unlike the AD equilibrium, the KKO model allows asymmetries in trade opportunities to arise from the underlying graph, yielding local price vectors that are generally distinct among agents.

Definition 2.3. A KKO equilibrium of a graphical economy is a pair (\mathbf{p}, \mathbf{x}) of prices $\mathbf{p} \in \mathbb{R}_+^{m \times \ell}$ and consumption plans $\mathbf{x} \in \mathbb{R}_+^{m \times m \times \ell}$, such that the following conditions are satisfied for all $i \in [m]$:

1. *Local Clearing.*

$$\sum_{j=i} \mathbf{x}^{ji} = \mathbf{e}^i. \quad (3)$$

2. *Individual rationality (IR).* Setting $\hat{\mathbf{x}}^i = \mathbf{x}^i$ maximizes the utility $u_i \left(\sum_{j=i} \hat{\mathbf{x}}^{ij} \right)$ over all $\hat{\mathbf{x}}^i \in \mathbb{R}_+^{m \times \ell}$ satisfying

$$\sum_{j=i} \mathbf{p}^j \cdot \hat{\mathbf{x}}^{ij} \leq \mathbf{p}^i \cdot \mathbf{e}^i. \quad (4)$$


As above, one can see eq. (3) as a clearing constraint by defining the demand vector from agent i by $\mathbf{d}^i(\mathbf{x}) = \sum_{j=i} \mathbf{x}^{ji}$. When G is the complete graph, we recover the AD equilibrium.

Graphical economies such as KKO are both technically and conceptually appealing. As is often the case in computer science, designing algorithms to exploit the underlying network topology can yield more efficient algorithms; see [22, 23, 31], as well as §5. Conceptually, as one of the most well-studied and general mathematical models of trade, placing AD on a network provides a natural way of understanding the effects of network topology in actual economies and, more generally, resource routing problems. Studying the relationship between network structure and *local* equilibrium outcomes yields better understanding of the phenomena observed in data [23].

Kakade et al. [22] show the existence of graphical equilibria under very general conditions, and give an algorithm to compute equilibria in polynomial time. Their algorithm follows from a novel algorithm to compute AD equilibria, together with a reduction to AD from the graphical setting by using $m\ell$ “tagged” goods which are marked according to who sold them. The model we present below will not allow such a reduction.

3 ALLOWING RESALE

Consider the following natural 3-node example: two agents with complementary endowments and preferences, each connected only to a single “broker” agent with no endowment. The utilities in this example are linear; here and throughout the paper when we work with linear utilities, we will specify the utility u_i by its coefficients $\mathbf{u}^i \in \mathbb{R}_+^\ell$, so that $u_i(\mathbf{v}) = \mathbf{u}^i \cdot \mathbf{v}$.



Intuitively, one might expect that agent 2 would extract rent from the other agents, as without a broker, agents 1 and 3 cannot trade and would have no utility. However, there is no KKO equilibrium, as technically agents 1 and 3 cannot trade directly even in the presence of agent 2. To see this, note that for the market to clear in eq. (3), no agent can purchase anything from agent 2, as the demand must match the supply $\mathbf{e}^2 = (0, 0)$. Furthermore, agents 1 or 3 cannot have any nonzero budget after selling their endowment, since otherwise they would rationally spend it on goods which they and their neighbors do not have in their endowments. We conclude that the prices of the nonzero endowments are zero: $p_1^1 = 0$ and $p_2^3 = 0$. But now that these prices are zero, agent 2 only satisfies individual rationality by buying an infinite amount of these goods, making clearing impossible. We conclude that no equilibrium exists under the KKO model.

A standard way to address this non-existence is to impose a small minimum endowment, i.e., to replace each zero in the endowment vectors with ϵ . Let us call an ϵ -KKO equilibrium one which arises from a KKO equilibrium after imposing this minimum; such an equilibrium always exists by [22]. In the ϵ -KKO equilibrium for this example, since agent 2 is indifferent between both goods, agent 2 extracts most of the utility as expected. In §3.2 we show that our formalism using resale leads to an equivalent outcome.

Unfortunately, ϵ -KKO equilibria do not always match intuition. Consider a ϵ -KKO equilibrium for a modification of the above example, where we set $\mathbf{u}^2 = (0, 1)$. Since agents 1 and 2 have no utility for the first good and one of them must consume this good for the market to clear in eq. (3), agent 1 must have zero price for that good: $p_1^1 = 0$. The endowment of agent 1 is therefore distributed among agents 1 and 2 for no profit and no utility, meaning an entire unit of utility is wasted in any ϵ -KKO equilibrium. By contrast, under our model agent 2 takes advantage of resale by facilitating the trade of the first good to agent 3. See Appendix C for a full analysis.

3.1 Resale Equilibrium

Motivated by the above examples, we introduce a new equilibrium concept for graphical economies which allows agents to facilitate trade between other parties without consuming any of the goods so traded; in other words, agents are allowed to *resell* goods. By introducing resale we are able to extend KKO equilibria to better capture the behaviour one expects in graphical economies at equilibrium. Though resale equilibria are defined as a one-shot process like the AD and KKO settings, it is intuitive to think of a resale equilibrium as proceeding in two phases: a resale phase and a purchase phase. In the resale phase, agents purchase goods from each other but immediately sell them, intuitively to arbitrage prices which differ within their neighborhood. At the end of the resale phase, the market need not clear. Next, in the purchase phase, agents sell their endowments, and then use this money, together with the profits from the first phase, to purchase goods to optimize their utility. Only after both phases do we require the market to clear at each node, in the sense

that every agent holds onto a nonnegative amount of each good, and no goods are created or destroyed in the economy.

To capture resale, we track purchases for resale separately from purchases for consumption. The vector $\mathbf{y}^{ij} \in \mathbb{R}_+^\ell$ denotes the bundle of goods agent $i \in [m]$ purchased from agent $j \in [m]$ for resale, and as before we define $\mathbf{y}^{ij} = \mathbf{0}$ for $j \neq i$. As with \mathbf{x}^i, \mathbf{x} , we define $\mathbf{y}^i = \{\mathbf{y}^{ij} : j \in [m]\}$ and $\mathbf{y} = \{\mathbf{y}^i : i \in [m]\}$ to be the resale plans for an agent and for the whole economy respectively. Finally, associated with each agent $i \in [m]$ is a resale bound $b_i \in \mathbb{R}_+$ representing an exogenous limit on resale for agent i . The vector $\mathbf{b} \in \mathbb{R}_+^m$ represents the set of every agent's resale bound. The resale limits in \mathbf{b} have a natural interpretation as credit or capacity constraints, as we describe below.

We define equilibria in graphical economies with resale in terms of *demand systems*. This formalism provides an intuitive means of reasoning about graphical economies with resale in their full generality and will be particularly important for the algorithm we introduce in §5. Our use of demand systems is inspired by the recent work of Garg et al. [13], though our algorithm differs from theirs in significant ways beyond just applying to a graphical setting with resale. Informally, demand systems encode optimal actions for agents in the economy at a given set of prices and budget. For example, in both AD and KKO equilibria, a demand system will return the set of consumption plans maximizing an agent's utility while satisfying eqs. 2 and 4 above.

Definition 3.1. A demand system is a function $D : \mathbb{R}_+^{m \times \ell} \times \mathbb{R}_+ \rightarrow 2^{\mathbb{R}_+^{m \times \ell}}$. A demand system is said to be *normalized* if $D(\mathbf{p}, 0) = \{\mathbf{0}\}$ for all $\mathbf{p} \in \mathbb{R}_+^{m \times \ell}$, and said to be *scale invariant* if $D(\mathbf{p}, \beta) = D(\alpha\mathbf{p}, \alpha\beta)$ for all $\alpha > 0, \beta \geq 0, \mathbf{p} \in \mathbb{R}_+^{m \times \ell}$.

In our model each agent i has two demand systems: a consumption demand system C_i and a resale demand system R_i . For consumption, $C_i(\mathbf{p}, \beta_i)$ denotes the set of plans \mathbf{x}^i that maximize agent i 's utility $u_i(\sum_{j \neq i} \mathbf{x}^{ij})$ at prices \mathbf{p} within the budget $\beta_i = \mathbf{p}^i \cdot \mathbf{e}^i + \sum_{j \neq i} (\mathbf{p}^i - \mathbf{p}^j) \cdot \mathbf{y}^{ij}$. Note that, as defined above, $\mathbf{x}^{ij} = \mathbf{0}$ for $j \neq i$ in every consumption plan $\mathbf{x}^i \in C_i(\mathbf{p}, \beta_i)$. When an agent has linear utilities \mathbf{u}^i , $C_i(\mathbf{p}, \beta_i)$ will include all consumption plans \mathbf{x}^i with fractional assignments of goods maximizing bang-per-buck u_k^i/p_k^j with total price β_i .²

For resale, agent i 's budget is b_i and $R_i(\mathbf{p}, b_i)$ intuitively denotes the set of resale plans \mathbf{y}^i that maximize resale profits $\sum_{j \neq i} (\mathbf{p}^i - \mathbf{p}^j) \cdot \mathbf{y}^{ij}$ at prices \mathbf{p} subject to some resale constraints expressed in terms of b_i . Recall that $\mathbf{y}^{ij} = \mathbf{0}$ for $j \neq i$ in every $\mathbf{y}^i \in R_i(\mathbf{p}, b_i)$. While our results apply to quite general resale demand systems, a natural and motivating example is *credit bound resale*, where b_i represents the total budget or "line of credit" with which agent i can purchase goods in order to resell. Formally, the credit bound resale demand system $R_i(\mathbf{p}, b_i)$ returns the set of plans \mathbf{y}^i maximizing $\sum_{j \neq i} (\mathbf{p}^i - \mathbf{p}^j) \cdot \mathbf{y}^{ij}$ over all \mathbf{y}^i satisfying $\sum_{j \neq i} \mathbf{p}^j \cdot \mathbf{y}^{ij} \leq b_i$, namely any fractional assignment of goods of total cost b_i maximizing profit-per-credit $p_k^i/p_k^j - 1$. Graphical economies with resale where

²When the optimal consumption plans are unique for each set of prices and budget, these systems coincide with demand functions, which are well studied in exchange economies for certain classes of utility functions. To see the importance of considering demand systems rather than demand functions, simply note that linear utility functions often have non-unique demand sets for a given set of prices and budgets.

every agent $i \in [m]$ has linear utilities and credit bound resale will be a running example throughout this paper.


Definition 3.2. Let C_i and R_i be the consumption and normalized resale demand systems of each agent $i \in [m]$. For any $\mathbf{b} \in \mathbb{R}_+^m$, a *b-resale equilibrium* of a graphical economy is a triple $(\mathbf{p}, \mathbf{x}, \mathbf{y})$ of prices \mathbf{p} , consumption plans \mathbf{x} , and resale plans \mathbf{y} , such that for all $i \in [m]$ the following conditions are satisfied:

1. *Local Clearing.* $\sum_{j \neq i} \mathbf{x}^{ji} + \sum_{j \neq i} \mathbf{y}^{ji} = \mathbf{e}^i + \sum_{j \neq i} \mathbf{y}^{ij}$.
2. *Optimal arbitrage.* $\mathbf{y}^i \in R_i(\mathbf{p}, b_i)$.
3. *Individual rationality (IR).* $\mathbf{x}^i \in C_i(\mathbf{p}, \mathbf{p}^i \cdot \mathbf{e}^i + \sum_{j \neq i} (\mathbf{p}^i - \mathbf{p}^j) \cdot \mathbf{y}^{ij})$.

For credit bound resale and similar resale demand systems, we recover the KKO equilibrium concept for $\mathbf{b} = \mathbf{0}$ and the AD equilibrium for sufficiently large \mathbf{b} . To see this, for every $i \in [m]$, consider a normalized resale system R_i that returns sets \mathbf{y}^i that are increasing in b_i when i can profit. To recover KKO, we simply set $\mathbf{b} = \mathbf{0}$ to remove resale: $\mathbf{b} = \mathbf{0}$ implies $\mathbf{y}^i = \mathbf{0}$ for all i as R_i is normalized, and thus consumption budgets are $\mathbf{p}^i \cdot \mathbf{e}^i$ which is equivalent to eq. (4) and market clearing simply reduces to eq. (3). Conversely, let b_i be sufficiently large for each agent. As resale is essentially unrestricted, then any difference in prices will result in large reselling. In particular, for all i , if b_i is large enough³ for there to exist $\mathbf{y}^i \in R_i(\mathbf{p}, b_i)$ at equilibrium prices \mathbf{p} such that \mathbf{y}^i includes every good in the economy, then the market will not clear if any prices differ. Thus, for sufficiently large \mathbf{b} , we will have $\mathbf{p}^i = \mathbf{p}^j$ for all i, j , forcing zero profits from resale. As agents are therefore indifferent among all valid resale plans, goods can move freely throughout the graph, and we can set \mathbf{y} to achieve any consumption plan \mathbf{x} satisfying global market clearing, eq. (1).

3.2 Revisiting the Broker Example

We can now see how the addition of resale changes the outcome of the broker example from above. For simplicity let R_i be a credit bound resale demand system for every agent $i \in [3]$. We will argue that, for any $0 < b \leq 2$, the following prices at $\alpha = \sqrt{b/2}$ lead to a *b-resale equilibrium* where $\mathbf{b} = (b, b, b)$. Dotted lines depict movement of the goods.



Let us verify the equilibrium. Agents 1 and 3 have nothing to gain from resale, so abstain from the first phase, whereas agent 2 can resell a total of b/α units. Agent 2 thus purchases $(b/2\alpha, 0)$ and $(0, b/2\alpha)$ from agents 1 and 3, respectively, to then resell for an optimal profit of $(1 - \alpha)b/\alpha$. In the second phase, agents 1 and 3 sell

³For credit bound resale, it suffices to set $b_i \geq \sum_{j \in [m], k \in [\ell]} p_k^j e_k^i$, the total cost of buying all goods in the economy.

their endowments for a revenue of α and purchase $(0, \alpha)$ and $(\alpha, 0)$ from agent 2, respectively, thus optimizing their utility. Agent 2 uses the profit from the first phase to purchase $(b(1 - \alpha)/2\alpha^2, 0)$ and $(0, b(1 - \alpha)/2\alpha^2)$ from agents 1 and 3, respectively, at a price α . Optimal arbitrage and individual rationality are clear; it remains to check the clearing constraint. Setting $\alpha = \sqrt{b/2}$, we now have that the goods consumed by agent 2 are $(1 - \alpha, 1 - \alpha)$, and the goods resold are indeed consumed by agents 1 and 3.

The resolution of the broker example is intuitive: adding a small amount of resale capacity b allows agent 2 to extract nearly all of the rent for the service of facilitating trade between agents 1 and 3. Specifically, the final utilities are $\sqrt{b/2}$, $2(1 - \sqrt{b/2})$, and $\sqrt{b/2}$, respectively. The larger b becomes, the less rent agent 2 can extract, for the simple reason that the prices for the other agents must increase for the market to clear. When $b \geq 2$, all prices become equal and the market is effectively a classic AD exchange economy as predicted in §3.1.

Finally, while this example is motivating as it yields no KKO equilibrium, and we must have $b > 0$ for the same reason, we can easily take a limit as $b \rightarrow 0$ to obtain a limiting equilibrium where agent 2 extracts *all* the rents from its neighbors. This leads to an intuitive extension of the KKO equilibrium concept, which remains well-defined even for agents with zero endowments: a limiting equilibrium of a b -resale economy as $b \rightarrow 0$. This limiting resale equilibrium yields an efficient allocation in both this example and the modification above with $\mathbf{u}^2 = (0, 1)$. In contrast, the corresponding limiting ϵ -KKO equilibrium as $\epsilon \rightarrow 0$ suffers from the same problem as the non-limit version: for the modified example with $\mathbf{u}^2 = (0, 1)$, the first good is “trapped” in a local economy regardless of ϵ , and thus in the limit.

4 EXISTENCE OF EQUILIBRIA

We prove the existence of equilibria in graphical economies with resale that satisfy five assumptions. Our assumptions are weaker than those by Arrow and Debreu [1] and Kakade et al. [22], and are closer to those made by Maxfield [27] and McKenzie [30], which are some of the mildest assumptions needed for proving the existence of equilibria in AD economies. Other than Assumptions 1 and 2 which are specific to graphical settings and resale, our assumptions are standard utility and non-degeneracy assumptions in line with those made by Arrow and Debreu [1], Kakade et al. [22], and Maxfield [27]. The assumptions specific to our setting are rather weak, and essentially state that the underlying graph G does not create pathological local economies. A formal derivation of our proof and the standard utility and non-degeneracy assumptions can be found in Appendix A.

To formally state Assumption 1, call a path between agents $i, j \in [m]$ a *trade path* if every node $\hat{i} \notin \{i, j\}$ on the path has $b_{\hat{i}} > 0$. We write $P(i, j)$ for the set of agents on some trade path between i and j ; by convention $i, j \in P(i, j)$, and by definition, $P(i, j) = P(j, i)$. Also, define the *supply graph of k* as the directed graph $G_k^S = ([m], E_k^S)$ with $(i, j) \in E_k^S$ if $e_k^i > 0$, j is non-satiable in k , and there is a trade path between i and j . Edges (i, j) in G_k^S are present whenever i could directly or indirectly sell their endowment of k to j , when price conditions allow for profitable resale in between, if applicable. Finally, let $G^S = ([m], E^S)$ be the (non-disjoint) union of supply

graphs on all goods, with $E^S = \bigcup_{k \in [\ell]} E_k^S$. In graph theoretic terms, G^S encodes all possible directed source-sink pairs in the economy, and each G_k^S encodes source-sink pairs for the good k .

Assumption 1. For every agent $i \in [m]$ and good $k \in [\ell]$:

- (i) If i is non-satiable in k then there exists an incoming edge to i in G_k^S .
- (ii) If $\mathbf{e}^i = \mathbf{0}$, there exists an edge in G_k^S from \hat{i} to \hat{j} such that $i \in P(\hat{i}, \hat{j})$ and $\mathbf{e}^{\hat{j}} \neq \mathbf{0}$.
- (iii) Furthermore, for each connected component C of the underlying economy graph G , the subgraph of G^S induced by $\{j \in C : \mathbf{e}^j \neq \mathbf{0}\}$ is strongly connected.

One technical interpretation of Assumption 1 is that we cannot partition any connected component of the underlying trade graph G into two sets of endowed agents where one group cannot supply any good the other group wants. This assumption is very closely related to the irreducibility assumption used to prove existence of AD equilibria by McKenzie [28, 29]. While McKenzie [28] introduces irreducibility on the entire economy as a sufficient assumption for proving existence, our assumption need only hold on each connected component. The supply graph G^S serves a similar purpose to the directed graphs constructed in Maxfield [27] for proving existence, though we place weaker assumptions on consumers in exchange for slightly stronger assumptions on G^S .

Our next assumption asserts that every agent in the economy must be able to participate with every good in the economy.

Assumption 2. For each consumer $i \in [m]$, $b_i > 0$ or $\mathbf{e}^i > \mathbf{0}$.

When $\mathbf{b} = \mathbf{0}$, Assumption 2 is analogous to assumptions made in both Arrow and Debreu [1] and Kakade et al. [22]. When $\mathbf{b} > \mathbf{0}$, every agent has some capacity to resell, and for sufficiently small \mathbf{b} , Assumption 2 behaves like a weaker form of its counterpart in those works.⁴ This assumption can be thought of as a near-minimal requirement to participate in the economy. It may be possible to replace $\mathbf{e}^i > \mathbf{0}$ with $\mathbf{e}^i \neq \mathbf{0}$, but our current proof techniques would require strengthening Assumption 1 for technical reasons related to the propagation of non-zero prices.

Assumptions 3, 4, and 5 are standard utility and non-degeneracy assumptions, in line with those made in [1, 22, 27]; we defer the details to Appendix A and summarize them here. Assumption 3 ensures that every good appears in some agent’s endowment, as otherwise the good can be ignored. Assumption 4 gives standard regularity conditions on the utility functions: continuity, non-satiability, and quasi-concavity. Assumption 5 states that resale works very loosely like credit bound resale: agents arbitrage prices if $b_i > 0$, and are not satiated in the presence of an infinite arbitrage (goods available for free to resell at some positive price).

THEOREM 4.1. In any graphical economy satisfying Assumptions 1–5, there exists a \mathbf{b} -resale equilibrium.

The conditions of Theorem 4.1 are straightforward to check. Given the utility functions, resale constraints, endowments, and underlying trade graph G in standard forms (e.g. encoded in matrices),

⁴Both Arrow and Debreu [1] and Kakade et al. [22] require that every agent is endowed with a non-zero amount of every good. When $\mathbf{b} > \mathbf{0}$, Assumption 2 allows agents to have sparse endowment vectors even when \mathbf{b} is set to such small values that resale is virtually non-existent in the economy.

we can construct G^S and check the conditions in Assumptions 1–5 in polynomial time using standard graph algorithms with additional bookkeeping. For example, consider checking the assumptions for the broker example in §3. Assumption 4 is easily verified for linear utilities; Assumption 5 is also easily verified for credit bound resale; Assumption 2 follows from all agents having $b_i > 0$; Assumption 3 is satisfied by the endowments of Agents 1 and 3; Assumption 1 is satisfied as Agents 1 and 3 are intermediated by Agent 2 and so all agents have access to the goods they want, Agent 2 is on every trade path, and G^S is connected on Agents 1 and 3.

5 COMPUTING EQUILIBRIA

Computing equilibria is known to be a computationally hard problem in general market models [16]. In addition to proving existence of equilibria in graphical economies with resale, we also introduce an auction-based algorithm that finds an approximate equilibrium. Our algorithm generalizes the results of Kapoor et al. [24] and Garg and Kapoor [18]. In this section we introduce the technical preliminaries needed for our algorithm along with its guarantees (§5.1). We then overview the algorithm (§5.2) and give proof sketches of correctness and time complexity (§5.3). Full pseudo-code and proofs are in Appendix B.

5.1 Preliminaries & Technical Guarantees

Two additional assumptions on demand systems are needed for our algorithm. These assumptions are standard in auction-based algorithms for approximating equilibria of markets [13, 18, 24]. First, we assume that the consumption and resale demand systems C_i and R_i are scale invariant for every agent $i \in [m]$, which is sufficient to guarantee that there exists an equilibrium for our algorithm to approximate for any initial prices. For example, if \mathbf{p} is a set of prices at a \mathbf{b} -resale equilibrium, then scale invariance ensures $\alpha\mathbf{p}$ is also a set of equilibrium prices for all $\alpha > 0$. The second assumption, weak gross substitutability, appears necessary for agents to myopically update consumption and resale plans without having to retract previous decisions.

Definition 5.1. We say that the demand system D satisfies the *weak gross substitutability* (WGS) property if for any $\hat{\mathbf{p}} \geq \mathbf{p}$, $\hat{\beta} \geq \beta$, and $\rho \in D(\mathbf{p}, \beta)$, there exists $\hat{\rho} \in D(\hat{\mathbf{p}}, \hat{\beta})$ such that $\hat{\rho}_k \geq \rho_k$ whenever $\hat{p}_k^j = p_k^j$ for $j \in [m]$ and $k \in [\ell]$. If D satisfies WGS, a *demand oracle* is an algorithm which produces such a $\hat{\rho}$ as output when given $\hat{\mathbf{p}}, \mathbf{p}, \hat{\beta}, \beta$, and $\rho \in D(\mathbf{p}, \beta)$ as input.

We assume that, for every agent $i \in [m]$, the consumption demand system C_i satisfies WGS and that the negation of the resale demand system $-R_i$ satisfies WGS.⁵ Intuitively, this assumption means that raising the price on a good does not decrease the consumption nor increase the resale of any other good. As such, consumption and resale demand systems that satisfy WGS ensure that local price raises do not leave agents regretting previous decisions. This property will be integral for the algorithm to find approximate equilibria efficiently since it will monotonically increase prices by regular increments. These assumptions are identical to those used

⁵By saying $-R_i$ satisfies WGS, we mean that the demand system resulting from negating every resale plan that is the output from R_i is WGS. From the reduction in [14], this is equivalent to satisfying WGS for production.

in Garg et al. [13], which introduces a different auction based algorithm for the pure exchange setting; we refer to that work for a comprehensive discussion of WGS consumption demand systems.

Throughout the algorithm, agents will consult a *consumption oracle* and a *resale oracle*, described in Definition 5.1. These oracles are algorithms that provide agents with a consumption and resale plan guaranteed to exist by the definitions of WGS demand systems. In the case of consumption, the oracle only cares about satisfying the WGS property of the consumption demand system. For resale, the oracles used in the algorithm also take into consideration an additional *request* vector of the form $\mathbf{r} \in \mathbb{R}_+^\ell$ as input, and try to have close to r_k of each good k . These resale oracles satisfy the WGS property of the negated resale demand system, and minimize $\|\mathbf{r} - \sum_{j \neq i} \hat{\mathbf{y}}^{ij}\|_1$ among all negated resale plans satisfying WGS. Taking requests helps resale oracles cater to the actual consumption needs of the economy when multiple resale plans maximize profits.

Our algorithm finds solutions that satisfy relaxed versions of the requirements for \mathbf{b} -resale equilibria (Definition 3.2). We use a notion of $(1 + \epsilon)$ -approximate \mathbf{b} -resale equilibrium that is consistent with Garg et al. [13], but uses the stronger notions of approximate clearing, and approximately efficient budget use, from Garg and Kapoor [18] and Kapoor et al. [24].

Definition 5.2. For an $\epsilon > 0$, the triple $(\mathbf{p}, \mathbf{x}, \mathbf{y})$ of prices \mathbf{p} , consumption plans \mathbf{x} , and resale plans \mathbf{y} is a $(1 + \epsilon)$ -approximate \mathbf{b} -resale equilibrium if for all $i \in [m]$ the following conditions are satisfied. Here $\bar{\mathbf{p}}$ satisfies $\bar{\mathbf{p}}^i = \mathbf{p}^i$ and $\bar{\mathbf{p}}^j = \mathbf{p}^j / (1 + \epsilon)$ for every $j \in [m] \setminus \{i\}$, and we define the budgets $\beta_i = \mathbf{p}^i \cdot \mathbf{e}^i + \sum_{j \neq i} (\mathbf{p}^i - \mathbf{p}^j) \cdot \mathbf{y}^{ij}$ and $\bar{\beta}_i = \bar{\mathbf{p}}^i \cdot \mathbf{e}^i + \sum_{j \neq i} (\bar{\mathbf{p}}^i - \bar{\mathbf{p}}^j) \cdot \mathbf{y}^{ij}$.

1. *Approx. Clearing.* $\frac{e^i + \sum_{j \neq i} y^{ij}}{(1 + \epsilon)} \leq \sum_{j \neq i} x^{ji} + \sum_{j \neq i} y^{ji} \leq e^i + \sum_{j \neq i} y^{ij}$.
2. *Approx. Optimal Arbitrage.* $\mathbf{y}^i \leq \bar{\mathbf{y}}^i$ for some $\bar{\mathbf{y}}^i \in R_i(\bar{\mathbf{p}}, b_i)$.
3. *Approx. Individual Rationality.* $\mathbf{x}^i \leq \bar{\mathbf{x}}^i$ for some $\bar{\mathbf{x}}^i \in C_i\left(\frac{\mathbf{p}}{1 + \epsilon}, \bar{\beta}_i\right)$.
4. *Approx. Efficient Budget Use.* $\frac{\beta_i}{1 + \epsilon} \leq \sum_{j \neq i} \mathbf{p}^j \cdot \mathbf{x}^{ij} \leq \beta_i$.

In other words, market clearing and budget spending are within a multiplicative factor of $(1 + \epsilon)$, while both consumption and resale plans are subsets of optimal plans within a multiplicative factor of $(1 + \epsilon)$ of prices and budgets.

Using the ideas introduced above, our algorithm achieves the following guarantees. As is often the case in auction-based algorithms, the time complexity is expressed in terms of the largest price found by the algorithm [13, 18, 24], which we denote by p_{max} . An upper bound on p_{max} can be found for specific sets of demand and resale systems, i.e. when these systems are given by explicit utility functions and resale constraints.

THEOREM 5.3. *Let $N_{max} = \max_{i \in [m]} |\{j \approx i\}|$ be the size of the largest neighborhood in G . For any graphical economy with resale satisfying Assumptions 1–5, and for which all consumption demand systems C_i and negated resale demand systems $-R_i$ are scale invariant and satisfy WGS, a $(1 + \epsilon)$ -approximate \mathbf{b} -resale equilibrium can be found in*

$$O\left(\frac{\ell}{\epsilon} C\left(mT_D + m^2 N_{max} T_R + m^3 \ell N_{max}^3\right) \log p_{max}\right)$$

time, where T_D is the time needed for one call to the demand oracle, T_R is the time needed for one call to the resale oracle, p_{max} is the largest price found by the algorithm, $M(R_i)$ is the maximal quantity of resold goods $\|\mathbf{y}^i\|_1$ for any $\mathbf{y}^i \in R_i(\mathbf{p}^*, b_i)$ and any $\mathbf{p}^* \in (0, p_{max}]^{m \times \ell}$, and $C \leq \frac{(1+\epsilon)p_{max}}{\epsilon e_{min}} \sum_{i \in [m], k \in [\ell]} (e_k^i + M(R_i))$.

5.2 Overview of an Auction-Based Algorithm for Graphical Economies

Our algorithm can be thought of as a series of localized auctions or negotiation processes, where agents simultaneously try to outbid one another for desired goods and utilize resale to leverage their position in the network (to maximize resale profits which they can use to place bids). Throughout the algorithm the consumption, resale, and prices of goods are modified in a way that ensures approximate individual rationality and approximately optimal arbitrage are always satisfied. As the process unfolds, relaxed versions of approximate clearing and approximately efficient budget use are gradually tightened until, at termination, the algorithm has converged on a $(1 + \epsilon)$ -approximate b -resale equilibrium. Loosely speaking each step in the algorithm can be broken into the following:

- Agents use the current prices in their neighborhood to query their consumption oracle. They use any surplus budget to bid on goods in the consumption oracle’s output plan that are not currently satisfied (e.g., due to neighboring agents previously being unable to match demand).
- Having received bids on goods, if an agent has insufficient endowment to satisfy demand, they use the current prices in their neighborhood and the received bids as resale requests to query the resale oracle. From the resale oracle’s output, they find any adjustments to their resale plan that will increase profits while better matching requests, they place bids appropriately by using their resale budget subject to the resale constraint(s).
- In response to the bidding processes from both consumption and resale, the prices on goods are monotonically increased one at a time when neighbors are over-demanding a good.

Throughout the course of the algorithm, agents *assign* goods to the highest bidders by renegeing on assignments made at lower bids—whenever a good’s price is raised, an agent stops receiving bids at the old price. Agents only ever increase prices on a good if they have renegeed on every assignment made at lower prices and are unable to profitably adjust their resale plans, as this indicates demand outweighs supply. Since prices only increase when *all* assignments at the old price are renegeed on and their good is being over-demanded at the current price, it follows that goods are only ever assigned at one of two prices: its current price and its old price. The fact that prices are monotonically increased, along with the WGS assumption, ensures that demand gradually decreases to match supply. The algorithm terminates when either the market locally clears exactly for every agent or every agent’s budget is nearly spent.

5.3 Overview of Algorithm’s Analysis

We now give proof sketches for the correctness and runtime of our algorithm (Theorem 5.3), with full proofs in Appendix B.2 and B.3.

Correctness. We begin with three observations that hold for the algorithm’s entire duration: (i) Prices are monotonically raised, so a good’s current price is always higher than its old price. (ii) Each agent’s budget is computed using the current prices in their neighborhood, i.e. assuming they are selling their entire endowment and optimally reselling at the current (higher) prices. (iii) Agents bid for resale and consumption according to their respective oracle and the current budget.

Since consumption and negated resale oracles satisfy WGS, from (i) we know that price raises can never increase plans suggested by an agent’s oracles. Therefore, from (iii), it follows that approximate individual rationality and approximately optimal arbitrage hold for every agent for the entirety of the algorithm. The right inequality in approximate clearing follows from the fact that assignments are never made with goods an agent does not have on hand. Similarly, the right inequality in approximately efficient budget use follows from (ii) and (iii) since budgets are always computed at the current prices and oracles always suggest plans within the budget. The left inequality in both approximate clearing and approximately efficient budget use follow from the way termination conditions are specified for the algorithm. When the algorithm terminates because all goods clear locally in the economy, these left inequalities hold trivially; for approximate clearing it is by definition, and for approximately efficient budget use because it implies no good is over-demanded and thus budgets are used fully. For the other termination condition, when every agent’s budget is nearly spent, showing that the left inequality holds requires some careful accounting of technical details about agents’ surplus budgets (Appendix B.2).

Run-Time. To analyze the time complexity of the algorithm, we bound two quantities: (i) the number of times prices can be raised, and (ii) the number of steps between price raises before the algorithm must terminate. Using p_{max} to derive (i) follows the same argument as many other auction-based algorithms [18, 24]. The argument notes that prices are never decreased, and prices are always increased by a factor of $(1 + \epsilon)$, and thus the number of price raises is bounded by $\ell \log_{1+\epsilon} p_{max}$. Our derivation of (ii) uses a method resembling an amortized analysis and differs considerably from previous work on auction-based algorithms. We observe that goods are never assigned at an old price, and therefore the total amount of budgets being spent on bids at old prices can only decrease between price raises. Thus, by using the total amount of budgets spent on goods at their old price as a potential, we derive a bound on the amount of time before this potential reaches zero. We then argue that when the potential is zero, no goods are being assigned at their old price in the entire economy, thus either the market exactly clears or some good is over demanded and a price must get raised. Putting these bounds together, and deriving the time complexity of each sub-routine in the algorithm, yields the time complexity in Theorem 5.3. Precise details of this analysis are given in Appendix B.3.


6 DISCUSSION

We have introduced a variant of the Kakade, Kearns, and Ortiz (KKO) [22] model of an exchange economy embedded on a graph, so that agents may propagate trade to distant parts of the graph by reselling goods. As motivated by our broker examples (§3.2 and

Appendix C), our model has several appealing properties. We discuss these properties further here, along with other considerations, extensions, and future directions.

6.1 Market Breadth

As we saw in the broker example (§3.2), the addition of resale allows goods to move more than one hop in the network. To underscore this point, consider a natural extension of the broker example, where agents having complementary endowments and preferences are connected by some arbitrarily long chain of brokers. For the same reason as before, no equilibrium exists under the KKO model, but we would hope that with resale the market is able to move goods through the underlying network so as to match demand in maximally distant “local markets”.



$$\begin{aligned} e^1 &= (1, 0) & e^2 &= (0, 0) & e^{m-1} &= (0, 0) & e^m &= (0, 1) \\ u^1 &= (0, 1) & u^3 &= (1, 1) & u^{m-1} &= (1, 1) & u^m &= (1, 0) \end{aligned}$$

As with the original broker example, it is simple to check that the conditions of Theorem 4.1 are satisfied when we set $b = (b, \dots, b)$ for any $b > 0$. Thus, there exists a b -resale equilibrium. We now argue that in any such equilibrium, some amount of good 1 must travel all the way from agent 1 to agent m , and good 2 from agent m to agent 1. As argued in §3, consider the case where agent 2 does not resell good 2 to agent 1. If agent 1 had any budget, the market would not clear, as agent 1 would request good 2 from herself or from agent 2. Thus, we must have $p_1^1 = 0$, but then there is no optimal consumption plan, as even with zero budget agent 2 would request an infinite amount of good 1 from agent 1. By symmetry, the same applies to agents m and $m - 1$. We conclude that, in any b -resale equilibrium in this example, goods must traverse the entire length of the network.

6.2 Exact Equilibria

In Section 5 we introduce an algorithm for finding *approximate* b -resale equilibria of graphical economies with resale. In general the computation of equilibria is a computationally difficult problem (see e.g. [16]). Garg and Kannan [14] introduce an algorithm for computing exact equilibria in AD production economies with polyhedral production sets, given the number of goods in the economy as a constant. Through a simple reduction of graphical economies with resale to a classic AD economy with production (as highlighted in the proof of Lemma A.2), we can see that the algorithm for computing exact equilibria in [14] directly applies to our setting but that the algorithm is only polynomial time when the number of goods and agents are both constant (due to the “tagging” of goods). That said, by simply replacing the global market clearing notion that is typically used in AD economies with our local notion of market clearing in all of their arguments, their algorithm for computing exact equilibria is recovered provided the number of goods is constant in the graphical economy.

6.3 Resale with Capacity Constraints

Along with the credit bound resale we used in examples throughout this paper, another natural form is *commodity bound resale*, where agents maximize resale profits subject to a hard constraint on the amount of goods being resold. Formally, resale demand systems R_i now return resale plans \mathbf{y}^i such that setting $\hat{\mathbf{y}}^i = \mathbf{y}^i$ maximizes $\sum_{j \neq i} (\mathbf{p}^i - \mathbf{p}^j) \cdot \hat{\mathbf{y}}^{ij}$ over all $\hat{\mathbf{y}}^i$ satisfying $\|\hat{\mathbf{y}}^i\|_1 \leq b_i$. Commodity bounds on resale can be thought of as modeling limitations on storage or bandwidth available to agents for reselling goods.

Graphical economies with commodity bound resale have many appealing properties: they are intuitive, they have straightforward analytic solutions to the broker and market breadth (§6.1) examples, the algorithm in §5 is simple to derive, etc. Unfortunately, since they do not satisfy Assumption 5(ii), graphical economies with commodity bound resale require much stronger assumptions in order to prove the general existence of equilibria and fall squarely outside the scope of most existence results found in theoretical economics (e.g. [1, 26–30, 33]).

6.4 Relevance to the Economic Literature

In order to address concerns about the centralized nature of traditional economic models, several recent works have introduced models of *decentralized* markets. These models often place markets on a multipartite graph, where edges represent two agents being capable of trade with one another, and the agents play a specific role in buyer-seller interactions possibly facilitated by intermediaries, e.g. [4, 6, 12, 19, 20]. Kakade et al. [22] introduced a generalization of the AD exchange model to a graphical setting, enabling a robust set of interactions between agents on an arbitrary set of goods which is lacking in other economic models on networks. Unfortunately, as we have argued, the model in [22] is “too local” and does not allow for any degree of intermediation between agents.

This paper generalizes the work in [22] to allow for robust intermediation in graphical economies. One could therefore consider our model as a unifying generalization of many models of trade in the literature of decentralized markets. Furthermore, by presenting two modes for equating graphical economies with AD economies,⁶ we equip the study of decentralized markets with a novel and mathematically rigorous means of understanding decentralized market models within the context of competitive AD equilibria.

6.5 Future Directions

The broker example in §3.2 illustrates the fact that a node can extract rent from their position in the network. Exploring a more precise connection between this advantage and graph-theoretic parameters such as degree centrality is an exciting direction for future work. It would also be interesting to explore the connections between our model and Arrow–Debreu with production.

⁶Kakade et al. [22] showed that AD exchange economies are equivalent to graphical economies with agents embedded on a complete graph. In §3 we argued that AD exchange economies are equivalent to graphical economies with a certain class of Resale demand systems and sufficiently large resale capacities. This gives two ways for graphical economies to interpolate between centralized and decentralized markets: (i) by changing the density of the underlying graph, and (ii) by changing resale capacities in the economy.

REFERENCES

- [1] Kenneth J Arrow and Gerard Debreu. 1954. Existence of an Equilibrium for a Competitive Economy. *Econometrica* 22, 3 (1954), 265–290.
- [2] V. Avasarala, H. Polavarapu, and T. Mullen. 2006. An Approximate Algorithm for Resource Allocation Using Combinatorial Auctions. In *2006 IEEE/WIC/ACM International Conference on Intelligent Agent Technology*. 571–578.
- [3] Dimitri P. Bertsekas. 1992. Auction algorithms for network flow problems: A tutorial introduction. *Computational Optimization and Applications* 1 (1992), 7–66.
- [4] Larry Blume, David Easley, Jon Kleinberg, and Eva Tardos. 2007. Trading Networks with Price-Setting Agents. In *Proceedings of the 8th ACM Conference on Electronic Commerce* (San Diego, California, USA) (*EC '07*). Association for Computing Machinery, New York, NY, USA, 143–151. <https://doi.org/10.1145/1250910.1250933>
- [5] Yann Bramoullé, Andrea Galeotti, and Brian Rogers. 2016. *The Oxford Handbook of the Economics of Networks*. Oxford University Press. <https://www.oxfordhandbooks.com/view/10.1093/oxfordhb/9780199948277.001.0001/oxfordhb-9780199948277>
- [6] Syngjoo Choi, Andrea Galeotti, and Sanjeev Goyal. 2017. Trading in Networks: Theory and Experiments. *Journal of the European Economic Association* 15, 4 (02 2017), 784–817. <https://doi.org/10.1093/jeaa/jvw016> arXiv:<https://academic.oup.com/jeaa/article-pdf/15/4/784/19518068/jvw016.pdf>
- [7] Bruno Codenotti, Benton McCune, Sriram Penumatcha, and Kasturi Varadarajan. 2005. Market Equilibrium for CES Exchange Economies: Existence, Multiplicity, and Computation.
- [8] C. Curescu and S. Nadjim-Tehrani. 2008. A Bidding Algorithm for Optimized Utility-Based Resource Allocation in Ad Hoc Networks. *IEEE Transactions on Mobile Computing* 7, 12 (2008), 1397–1414.
- [9] Gerard Debreu. 1962. New Concepts and Techniques for Equilibrium Analysis. *International Economic Review* 3, 3 (1962), 257–273. <http://www.jstor.org/stable/2525394>
- [10] Nikhil R. Devanur. 2004. The Spending Constraint Model for Market Equilibrium: Algorithmic, Existence and Uniqueness Results. In *Proceedings of the Thirty-Sixth Annual ACM Symposium on Theory of Computing* (Chicago, IL, USA) (*STOC '04*). Association for Computing Machinery, New York, NY, USA, 519–528. <https://doi.org/10.1145/1007352.1007431>
- [11] Nikhil R. Devanur, Christos H. Papadimitriou, Amin Saberi, and Vijay V. Vazirani. 2008. Market Equilibrium via a Primal–Dual Algorithm for a Convex Program. *J. ACM* 55, 5, Article 22 (Nov. 2008), 18 pages. <https://doi.org/10.1145/1411509.1411512>
- [12] Douglas M. Gale and Shachar Kariv. 2009. Trading in Networks: A Normal Form Game Experiment. *American Economic Journal: Microeconomics* 1, 2 (2009), 114–132. <http://www.jstor.org/stable/25760364>
- [13] Jugal Garg, Edin Husić, and László A. Végh. 2019. Auction Algorithms for Market Equilibrium with Weak Gross Substitute Demands. arXiv:1908.07948 [cs.GT]
- [14] Jugal Garg and Ravi Kannan. 2015. Markets with Production: A Polynomial Time Algorithm and a Reduction to Pure Exchange. In *Proceedings of the Sixteenth ACM Conference on Economics and Computation* (Portland, Oregon, USA) (*EC '15*). Association for Computing Machinery, New York, NY, USA, 733–749. <https://doi.org/10.1145/2764468.2764517>
- [15] Jugal Garg, Ruta Mehta, Vijay V. Vazirani, and Sadra Yazdanbod. 2014. Leontief Exchange Markets Can Solve Multivariate Polynomial Equations, Yielding FIXP and ETR Hardness. *CoRR abs/1411.5060* (2014). arXiv:1411.5060 <http://arxiv.org/abs/1411.5060>
- [16] Jugal Garg, Ruta Mehta, Vijay V. Vazirani, and Sadra Yazdanbod. 2017. Settling the Complexity of Leontief and PLC Exchange Markets under Exact and Approximate Equilibria. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing* (Montreal, Canada) (*STOC 2017*). Association for Computing Machinery, New York, NY, USA, 890–901. <https://doi.org/10.1145/3055399.3055474>
- [17] Jugal Garg and Vijay V. Vazirani. 2013. On Computability of Equilibria in Markets with Production. arXiv:1308.5272 [cs.GT]
- [18] Rahul Garg and Sanjiv Kapoor. 2006. Auction Algorithms for Market Equilibrium. *Math. Oper. Res.* 31, 4 (Nov. 2006), 714–729. <https://doi.org/10.1287/moor.1060.0216>
- [19] Lining Han and Ruben Juarez. 2018. Free intermediation in resource transmission. *Games and Economic Behavior* 111 (2018), 75 – 84. <https://doi.org/10.1016/j.geb.2018.06.006>
- [20] Toomas Hinnoaar. 2019. Price Setting on a Network. *CoRR abs/1904.06757* (2019). arXiv:1904.06757 <http://arxiv.org/abs/1904.06757>
- [21] Kamal Jain and Kasturi Varadarajan. 2006. Equilibria for economies with production: Constant-returns technologies and production planning constraints. In *Proceedings of the seventeenth annual ACM-SIAM symposium on Discrete algorithm*. Society for Industrial and Applied Mathematics, 688–697.
- [22] Sham M Kakade, Michael Kearns, and Luis E Ortiz. 2004. Graphical economics. In *International Conference on Computational Learning Theory*. Springer, 17–32.
- [23] Sham M. Kakade, Michael Kearns, Luis E. Ortiz, Robin Pemantle, and Siddharth Suri. 2004. Economic Properties of Social Networks. In *Proceedings of the 17th International Conference on Neural Information Processing Systems* (Vancouver, British Columbia, Canada) (*NIPS '04*). MIT Press, Cambridge, MA, USA, 633–640.
- [24] Sanjiv Kapoor, Aranyak Mehta, and Vijay Vazirani. 2005. An Auction-Based Market Equilibrium Algorithm for a Production Model. In *Proceedings of the First International Conference on Internet and Network Economics* (Hong Kong, China) (*WINE '05*). Springer-Verlag, Berlin, Heidelberg, 102–111. https://doi.org/10.1007/11600930_11
- [25] Chengyu Lin, Tuo Yu, Riheng Jia, Feng Yang, and Xiaoying Gan. 2013. Auction based channel allocation in multi-hop networks. *2013 International Conference on Wireless Communications and Signal Processing* (2013), 1–6.
- [26] Andreu Mas-Colell, Michael D Whinston, and Jerry R Green. 1995. *Microeconomic Theory*. Oxford University Press.
- [27] Robert R. Maxfield. 1997. General equilibrium and the theory of directed graphs. *Journal of Mathematical Economics* 27, 1 (February 1997), 23–51. <https://ideas.repec.org/a/eee/mateco/v27y1997i1p23-51.html>
- [28] Lionel W. McKenzie. 1959. On the Existence of General Equilibrium for a Competitive Market. *Econometrica* 27, 1 (1959), 54–71. <http://www.jstor.org/stable/1907777>
- [29] Lionel W. McKenzie. 1961. On the Existence of General Equilibrium: Some Corrections. *Econometrica (pre-1986)* 29, 2 (04 1961), 247.
- [30] Lionel W. McKenzie. 1981. The Classical Theorem on Existence of Competitive Equilibrium. *Econometrica* 49, 4 (1981), 819–841. <http://www.jstor.org/stable/1912505>
- [31] Noam Nisan, Tim Roughgarden, Eva Tardos, and Vijay V. Vazirani. 2007. *Algorithmic Game Theory*. Cambridge University Press, USA.
- [32] Vijay V. Vazirani. 2010. Spending Constraint Utilities with Applications to the Adwords Market. *Math. Oper. Res.* 35, 2 (May 2010), 458–478. <https://doi.org/10.1287/moor.1100.0450>
- [33] Léon Walras. 1874. *Éléments d'économie politique pure, ou, Théorie de la richesse sociale*. L. Corbaz Lausanne.
- [34] Michael M. Zavlanos, Leonid Spesivtsev, and George J. Pappas. 2008. A distributed auction algorithm for the assignment problem. *2008 47th IEEE Conference on Decision and Control* (2008), 1212–1217.

A FULL PROOF AND MISSING ASSUMPTIONS FROM §4

A.1 Missing Assumptions

In addition to Assumptions 1 and 2 introduced in § 4, three standard utility and non-degeneracy assumptions are needed to prove Theorem 4.1. The first such assumption simply guarantees that every good exists at least somewhere in the economy. Intuitively, this assumption asserts that in order for a good to influence a competitive economy it needs to exist. Together with Assumption 1 introduced below, this assumption ensures all goods in the economy play a role.

Assumption 3. For every good $k \in [\ell]$, there exists an agent $i \in [m]$ such that $e_k^i > 0$.

The next assumption is a set of conditions on agents’ utility functions, which are typical in mathematical economics and match those needed in [1, 22, 27].

Assumption 4. For all agents $i \in [m]$, prices \mathbf{p} , and $\mathbf{x}^i \in C_i(\mathbf{p}, \mathbf{p}^i \cdot \mathbf{e}^i + \sum_{j \neq i} (\mathbf{p}^i - \mathbf{p}^j) \cdot \mathbf{y}^{ij})$, the utility function u_i such that consumption plans $\mathbf{x}^i \in \arg\max u_i(\sum_{j \neq i} \mathbf{x}^{ij})$ satisfies:

- (i) (Continuity) u_i is a continuous function.
- (ii) (Non-Satiability) There exists $k \in [\ell]$ such that u_i is strictly monotone increasing in k .
- (iii) (Quasi-Concavity) Let $\mathbf{x}, \mathbf{x}' \in \mathbb{R}_+^\ell$. If $u_i(\mathbf{x}') > u_i(\mathbf{x})$ then $u_i(\alpha \mathbf{x}' + (1 - \alpha)\mathbf{x}) > u_i(\mathbf{x})$ for all $0 < \alpha \leq 1$.

The utility conditions of Assumption 4 in turn influence the agents’ consumption demand systems. As was true in [22], Assumption 4 in the graphical setting, together with individual rationality, implies that in any \mathbf{b} -resale equilibrium agents only demand a good

for consumption at the cheapest price available in their neighborhood. Stated formally, by Assumption 4 and individual rationality, for agents $i \in [m]$, $j \simeq i$, and good $k \in [\ell]$, if $x_k^{ij} > 0$ then $p_k^j \leq p_k^i$ for all $\hat{j} \simeq i$. In addition, non-satiability ensures every agent i has at least one good k that i demands an infinite amount of if offered a zero price; in discussions to come, when we say i is *non-satiable in k* we are referring to this phenomenon.

Assumption 5 establishes weak conditions on resale demand systems so that satisfying optimal arbitrage behaves as we might typically expect in economic models. Firstly the assumption ensures that if an agent can resell profitably, then they do resell something without discriminating against specific goods or agents as long as it is profitable. Secondly the assumption ensures that if an agent i can resell some good $k \in [\ell]$ profitably, and a neighbor is offering k for free, then R_i will be empty since no finite amount of good k can saturate i 's credit (analogous to non-satiation in consumption). Finally, the assumption ensures that agents without the capacity to resell ($b_i = 0$) do not resell.

Assumption 5. For all agents $i \in [m]$, R_i is a normalized resale demand system and, for $b_i > 0$, if there exists $j \sim i$, $k \in [\ell]$ such that:

- (i) $0 < p_k^j < p_k^i$, then for all $\mathbf{y}^i \in R_i(\mathbf{p}, b_i)$, $\|\mathbf{y}^i\|_1 > 0$.
- (ii) $0 = p_k^j < p_k^i$, then $R_i(\mathbf{p}, b_i)$ will be empty.

A.2 Proof of Theorem 4.1

Our proof follows a similar trajectory to [22], but requires substantially more work due to our relaxations of their assumptions; in particular, endowments may be zero, and as such desired goods may not be directly available in an agent's neighborhood. At a high level, the proof strategy can be thought of as a three step process: a proof that *quasi-equilibria* exist, a weaker equilibrium concept which relaxes rationality; showing these assumptions suffice for quasi-equilibria to be \mathbf{b} -resale equilibria when the underlying graph is connected; and a proof that these assumptions guarantee the general existence of \mathbf{b} -resale graphical equilibria by establishing equilibria in every connected component of the underlying graph.

The concept of quasi-equilibria was first introduced by Debreu [9] to do away with the assumption in [1] that every agent in the economy has non-zero endowments of every good. (In particular, as discussed in §3, these criticisms apply to ϵ -KKO equilibria.) McKenzie [30] gives a detailed discussion of quasi-equilibria and AD equilibria, while Maxfield [27] shows that analogous quasi-equilibria exist in AD economies with production. Intuitively, quasi-equilibria are relaxations of equilibria: agents with wealth behave exactly as they would at equilibrium, while agents without wealth need not satisfy individual rationality. As rationality of agents with zero wealth applies only to goods with zero prices, quasi-equilibria are used as a stepping stone to ultimately rule out zero prices and/or zero wealth, in which case they coincide with the usual equilibria.

We say an agent i is *rational* if optimal arbitrage and individual rationality (conditions 2 and 3 of Definition 3.2) are both satisfied. Our definition of \mathbf{b} -resale quasi-equilibria is equivalent to its counterpart in AD production economies.

Definition A.1. A \mathbf{b} -resale quasi-equilibrium of a graphical economy with resale is a set of globally normalized prices \mathbf{p} (i.e. $\sum_{i \in [m], k \in [\ell]} p_k^i < \infty$), a set of consumption plans \mathbf{x} , and a set of resale plans \mathbf{y} , in which

the local markets clear, optimal arbitrage is enforced, and for each agent i with wealth $\beta_i = \mathbf{p}^i \cdot \mathbf{e}^i + \sum_{j \simeq i} (\mathbf{p}^i - \mathbf{p}^j) \cdot \mathbf{y}^{ij}$, the following conditions hold:

1. (Rational) If agent i has positive wealth ($\beta_i > 0$), then $\mathbf{x}^i \in C_i(\mathbf{p}, \beta_i)$, and thus i is rational.
2. (Quasi-Rational) If agent i has no wealth ($\beta_i = 0$), then \mathbf{x}^i is budget constrained, but need not be in $C_i(\mathbf{p}, \beta_i)$, i.e., individual rationality need not hold.

LEMMA A.2. In any graphical economy with a resale in which Assumptions 4 and 5 hold, there exists a \mathbf{b} -resale quasi-equilibrium.

PROOF. Our proof is straightforward—we satisfy all the preconditions for quasi-equilibrium existence highlighted in [27]. Notice that a graphical economy with resale is an AD production economy with $m\ell$ “tagged goods” and a “production firm” per agent who can resell. Let (i, k) for $i \in [m]$ and $k \in [\ell]$ represent the corresponding “tagged good” in the AD setting. Then the consumption plans in the AD setting only allow consumer i to consume goods (j, k) for $j \simeq i$. Similarly, each $i \in [m]$ such that $b_i > 0$ has sole ownership of a firm whose resale plans only allow input of goods (j, k) for $j \sim i$ and, for each input (j, k) , outputs an equal amount of (i, k) where the allowable consumption set is defined by i 's resale bounds. In this AD setting and with our assumptions it is clear that: the consumption sets are closed convex sets bounded from below, the production sets are closed convex sets containing $\mathbf{0}$ that do not produce goods “from nothing”, the utility functions still satisfy Assumption 4, and the economy satisfies “free disposal”. As stated in Maxfield [27] and McKenzie [30], these conditions are sufficient to imply the existence of a quasi-equilibrium in an AD production economy. \square

As noted above, quasi-equilibria are a powerful tool for proving existence, for the following reasons. First, agents can only consume while being quasi-rational if some good in their neighborhood has price zero since they have no wealth. Second, if all agents are behaving rationally, then the \mathbf{b} -resale quasi-equilibrium is exactly a \mathbf{b} -resale equilibrium. These facts play a crucial role in our proof—as they did in e.g. [9, 22, 27, 30]—because once we know a quasi-equilibrium exists, it suffices to show that all endowed agents are rational and all non-endowed agents have strictly positive local price vectors. Given these statements, an agent either has positive wealth, and thus is rational, or is forced to consume nothing.

With quasi-equilibria in hand, we turn to assumptions that ensure all agents are rational. Together Assumptions 1–5 provide powerful guarantees: certain non-zero prices must propagate in the presence of a rational agent (Lemma A.3), and any \mathbf{b} -resale quasi-equilibrium must have at least one *endowed* rational agent (Lemma A.4).

LEMMA A.3. Suppose we have a graphical economy with resale satisfying Assumptions 5 and 2. Let $i \in [m]$ be a rational agent, and suppose there exists a trade path between i and $j \in [m]$. For any \mathbf{b} -resale quasi-equilibrium $(\mathbf{x}, \mathbf{y}, \mathbf{p})$, and any $k \in [\ell]$, either of the following conditions imply $p_k^j > 0$ for all $\hat{j} \in P(i, j)$:

- (i) i is non-satiable in k , or
- (ii) $b_i > 0$ and $p_k^i > 0$.

PROOF. Let $i, j \in [m]$ be agents such that i is rational and there exists a trade path between i and j . If i is non-satiated in k , then by definition $p_k^{j'} > 0$ for all $j' \simeq i$; in particular, $p_k^i > 0$. Similarly, if $b_i > 0$ and $p_k^i > 0$, then $p_k^{j'} > 0$ for all $j' \simeq i$, as otherwise Assumption 5(ii) contradicts optimal arbitrage.

Arguing by contradiction, suppose $p_k^{\bar{j}} = 0$ for some $\bar{j} \in P(i, j)$. Let π be a trade path between i and j containing \bar{j} , ordered from i to j , and let (i', j') be the first edge along π such that $p_k^{i'} > 0$ and $p_k^{j'} = 0$. (Such an edge must exist as $p_k^i > 0$ and $p_k^{\bar{j}} = 0$.) Note that $i' \neq i$ as otherwise $j' \simeq i$ but $p_k^{j'} = 0$, contradicting the above; similarly, $i' \neq j$ as then $j' = j$ as well and p_k^j cannot be both positive and zero. Thus, $b_{i'} > 0$, as i' is on a trade path between i and j but $i' \notin \{i, j\}$, and by Assumption 5(ii) we therefore contradict optimal arbitrage. \square

LEMMA A.4. *If a connected graphical economy with resale satisfies Assumptions 4, 5, 2, 1, and 3, then for any \mathbf{b} -resale quasi-equilibrium there exists some rational agent $i \in [m]$ such that $\mathbf{p}^i \cdot \mathbf{e}^i > 0$.*

PROOF. Let \mathbf{x}, \mathbf{y} , and \mathbf{p} be a \mathbf{b} -resale quasi-equilibrium. By price normalization there exists at least one agent that has a good with non-zero price, i.e. there exists $i \in [m]$ such that $\mathbf{p}^i \neq \mathbf{0}$. Consider i 's wealth $\beta_i = \mathbf{p}^i \cdot \mathbf{e}^i + \sum_{j \simeq i} (\mathbf{p}^i - \mathbf{p}^j) \cdot \mathbf{y}^{ij}$. Note that, by definition, $0 \leq \sum_{j \simeq i} (\mathbf{p}^i - \mathbf{p}^j) \cdot \mathbf{y}^{ij}$. It follows that if $\mathbf{p}^i \cdot \mathbf{e}^i > 0$, then we know i has positive wealth and so i is rational. On the other hand, if $\mathbf{p}^i \cdot \mathbf{e}^i = 0$, then by Assumption 2 we know that $b_i > 0$.

Let $k \in [\ell]$ be a good such that $p_k^i > 0$. We now show that there exists a trade path between i and some agent $j \in [m]$ with $e_k^j > 0$. Note that some $j \in [m]$ with $e_k^j > 0$ exists by Assumption 3. By non-satiability (Assumption 4(ii)), agent i is non-satiable in some good $\hat{k} \in [\ell]$. By Assumption 1(i) for i and \hat{k} , there exists a trade path between i and some $\hat{j} \in [m]$ with $e_{\hat{k}}^{\hat{j}} > 0$. Furthermore, as $\mathbf{e}^{\hat{j}} \neq \mathbf{0}$, by Assumption 1(iii) there is a directed path $j_1 = \hat{j}, j_2, \dots, j_H = j$ in G^S from \hat{j} to j . For each $h \in [H]$, by Assumption 2, either $e_k^{j_h} > 0$ or $b_{j_h} > 0$. Let $h^* \in [H]$ be the first index such that $e_k^{j_{h^*}} > 0$, which must exist as $j = j_H$. We now have a trade path between i and $\hat{j} = j_1$, and j_h and j_{h+1} for all $1 \leq h < h^*$. As $b_{j_h} > 0$ for all $h < h^*$, stitching these trade paths together gives a trade path π between i and j_{h^*} . (Observe that if $H = 1$, many of the above statements are vacuously true, and the statement follows as $j_{h^*} = j = \hat{j}$.) We have established $j_{h^*} \in P(i, j_{h^*})$, and Lemma A.3 now gives $p_k^{j_{h^*}} > 0$. Therefore $\mathbf{e}^{j_{h^*}} \cdot \mathbf{p}^{j_{h^*}} > 0$, so j_{h^*} is rational. \square

Using similar arguments to those used to prove Lemmas A.3 and A.4, we find that rationality propagates in connected components of graphical economies much like non-zero prices do.

LEMMA A.5. *If a connected graphical economy with resale satisfies Assumptions 1–5, then for any \mathbf{b} -resale quasi-equilibrium it holds that every agent is rational.*

PROOF. Our proof proceeds in two stages. First we show that if there exists a rational endowed agent, then that rationality will “propagate” to all other endowed agents. We conclude by showing that this implies every non-endowed agents in the economy has

strictly positive local price vectors, which forces all non-endowed agents to behave rationally at a \mathbf{b} -resale quasi-equilibrium too.

Let \mathbf{x}, \mathbf{y} , and \mathbf{p} be a \mathbf{b} -resale quasi-equilibrium as guaranteed by Lemma A.2 and Assumptions 4 and 5. Suppose i is rational and non-satiated in k . For any edge (j, i) in G_k^S , we have $e_k^j > 0$ and there exists a trade path between i and j . By Lemma A.3, we have $p_k^{\hat{j}} > 0$ for all $\hat{j} \in P(i, j)$, and as $j \in P(i, j)$, we have $\mathbf{p}^j \cdot \mathbf{e}^j > 0$ and j is rational. By definition of G^S , if any agent i is rational, then all agents with a directed path to i in G^S are therefore rational, by repeating this argument for the respective goods on each edge of this path. From Lemma A.4, there exists a rational agent $i \in [m]$ with $\mathbf{e}^i \neq \mathbf{0}$. By Assumption 1(iii), G^S is strongly connected among endowed agents, so all agents j with $\mathbf{e}^j \neq \mathbf{0}$ have directed paths in G^S to i , and are therefore rational.

It remains to establish rationality of unendowed agents. By Assumption 1(ii), if $\mathbf{e}^j = \mathbf{0}$, then for all $k \in [\ell]$ there exists an edge (\hat{i}, \hat{j}) in G_k^S such that $\mathbf{e}^{\hat{j}} \neq \mathbf{0}$ and $j \in P(\hat{i}, \hat{j})$. By definition of G_k^S , we also have $\mathbf{e}^{\hat{i}} \neq \mathbf{0}$. By the above use of Lemma A.3, as $j \in P(\hat{i}, \hat{j})$, we must have $p_k^j > 0$. Therefore, for all unendowed j and goods k , $p_k^j > 0$. Letting $\beta_j = \mathbf{p}^j \cdot \mathbf{e}^j + \sum_{i \simeq j} (\mathbf{p}^i - \mathbf{p}^j) \cdot \mathbf{y}^{ij}$ be the budget for agent j , we conclude that if $\mathbf{e}^j = \mathbf{0}$, either $\beta_j = 0$, in which case $C_j(\mathbf{p}, \beta_j) = \{\mathbf{0}\}$ by Assumption 4 and therefore j rationally consumes nothing, or $\beta_j > 0$, in which case j is also rational. \square

Combining all five assumptions, and Lemma A.5, we are now ready to show existence of equilibria (Theorem 4.1). It remains only to show is that Lemma A.5 holds on each connected component of a graphical economy. By definition, if Assumptions 4, 5, 2, and 1 hold for a disconnected graphical economy with resale, then the assumptions also hold for each connected component of the graphical economy. In addition, if Assumption 3 does not hold for some connected component C of the underlying trade graph G , then there is some $k \in [\ell]$ with $e_k^i = 0$ for every $i \in C$. As Assumption 1 asserts that each i has an incoming edge in G_k^S , for each $k' \in [\ell]$ that i is non-satiable in, and that C is strongly connected among suppliers in G^S , it follows that there cannot exist an agent $j \in C$ that is non-satiable in k . Therefore, by appropriately scaling their prices, we know that goods contradicting Assumption 3 can be neglected at equilibrium in C , and Assumptions 1–5 hold in C for the set of goods $[\ell] \setminus \{k \in [\ell] : e_k^i = 0 \text{ for every } i \in C\}$.

As was just discussed, each connected component C of the graphical economy's underlying trade graph G satisfies Assumptions 1–5 (possibly by removing some “useless goods”). Thus, by Lemma A.5, every agent in each connected component C of the graphical economy is rational and, by definition, C has its own \mathbf{b} -resale equilibrium. Finally, by definition, the union of equilibria in uncoupled graphical economies is an equilibrium of the entire economy and so we know that a \mathbf{b} -resale equilibrium exists in any graphical economy satisfying our five assumptions.

B DETAILED AUCTION-BASED ALGORITHM AND ITS ANALYSIS

A sketch of the auction-based algorithm for computing approximate resale equilibria was given in §5, but several technical details were omitted for the sake of clarity and exposition—we will now expand

on all of these details. A detailed description of the algorithm is given in Appendix B.1, with complete pseudo-code given in Algorithm 1. Appendix B.2 and B.3 contain full proofs of correctness and run-time of the algorithm respectively, together these formally prove Theorem 5.3 as was sketched in §5.3.

Throughout Algorithm 1 we will use $\mathcal{O}_{C_i}(\hat{\mathbf{p}}, \hat{\beta}_i, \mathbf{x}^i)$ as shorthand meaning “agent i consults their consumption oracle for C_i with the input $(\mathbf{p}, \hat{\mathbf{p}}, \beta_i, \hat{\beta}_i, \mathbf{x}^i)$ ” where \mathbf{p} and β_i are the prices and budgets used during the previous call to the oracle. Similarly, we use $\mathcal{O}_{R_j}(\hat{\mathbf{p}}, b_j, \mathbf{y}^j, \mathbf{r})$ as shorthand for “agent j consults their resale oracle for R_j with the input $(\mathbf{p}, \hat{\mathbf{p}}, b_j, \mathbf{y}^j)$ and the request vector \mathbf{r} ” where \mathbf{r} is the “request” vector discussed in section §5.1 and \mathbf{p} is the set of prices used during the previous call to the oracle. Since prices are never decreased in Algorithm 1, maintaining prices and budgets from previous oracle calls involves straightforward bookkeeping.

B.1 Auction-Based Algorithm for Computing Approximate Equilibria

At initialization of the algorithm, a small accuracy parameter $\epsilon > 0$ is appropriately chosen, all prices in \mathbf{p} are set to unity, all consumption plans \mathbf{x} to zero, and all resale plans \mathbf{y} to zero. Prices will only ever increase throughout the algorithm and these increases are in multiplicative increments of $(1 + \epsilon)$. At the initial prices no agent can profit from resale, so only endowed agents have surplus budgets at the start of the algorithm. As mentioned in §5.2, a key concept in our discussion of the algorithm is that of goods being *assigned* at a price. Assignment simply refers to an agent “committing” a portion of a good they have available (from their endowment or purchased in order to resell) to an agent who has placed a bid on the good. Note that since both consumption and resale begin at zero, we say that all goods in the economy are initially unassigned. As will be discussed below, the algorithm assigns goods to the highest bidder and goods are only ever unassigned at initialization or when an agent gets outbid.

After initialization Algorithm 1 proceeds by iterating through the main loop until one of two conditions are met: the market locally clears exactly for every agent, or every agent’s budget is nearly spent. The amount of money in an agent’s budget that is not being used (i.e. not used in bids) is called that agent’s *surplus budget*. Within the main loop of Algorithm 1, agents are considered one by one. If some agent i is considered without surplus budget, then the algorithm simply moves onto the next agent. If some agent i is considered and does have positive surplus budget, then they consult their demand oracle and use their surplus budget to place bids according to the oracle’s output consumption plan. Suppose agent i has placed a bid on agent j ’s good k according to the demand oracle’s output plan. Agent j will try to meet the demand of i ’s bid via calls to the Assign, Outbid, and Reschedule_Resale procedures in that order.

In the Assign procedure agent j checks if it has some unassigned amount of good k and, if so, simply gives agent i an appropriate amount (either as much as j has or as much as i ’s bid can pay for) of agent j ’s good k at the current price p_k^j . If agent j was unable to meet agent i ’s demand for good k after Assign is called, then all of j ’s good k is already assigned to other agents at either the current

price (p_k^j) or the old price ($\frac{p_k^j}{(1+\epsilon)}$) and the Outbid procedure is called. In the Outbid procedure agent j “takes back” good k from agents who were assigned k at the old price ($\frac{p_k^j}{(1+\epsilon)}$), returns the money those agents had spent on bids for k at the old price, and reassigns the newly available quantities of k to agent i at the new price p_k^j . Finally, if agent j was unable to meet agent i ’s demand for good k after Outbid is called, then all of j ’s good k is assigned at the current price p_k^j and j will need to consider resale as a means of meeting agent i ’s demand for k ; this is when the Reschedule_Resale procedure is called.

In the Reschedule_Resale procedure agent j consults their resale oracle and, if the oracle has deemed it profitable, uses the oracle’s output to try and acquire more of good k for agent i by using credit. Though perhaps mysterious at first glance, calls to Reschedule_Resale essentially boil down to agent j comparing the resale oracle’s previous output with the current output and, if there is a profitable way to (re)allocate the credit being used, agent j tries procuring more of good k by using credit to place bids via internal calls to the Assign, Outbid, and Reschedule_Resale procedures. It is important that, by definition, j ’s resale oracle will never reduce resale profits. If after calling Reschedule_Resale agent j is still unable to satisfy agent i ’s bids for good k , then j knows it has priced k too low and will raise p_k^j by a fixed multiplicative factor of $(1 + \epsilon)$. At this point agent j updates its budget and the current iteration of the main loop is finished.

As alluded to above, at any given point in the algorithm goods can only be assigned at either the current price p_k^j or the previous price ($\frac{p_k^j}{(1+\epsilon)}$). This fact follows immediately from prices only being raised when a good is fully assigned at the current price and still over demanded (i.e. after a call to Reschedule_Resale). Furthermore, notice that the Reschedule_Resale procedure tries to procure goods by essentially mirroring the process in the main loop (which is key to our analysis in Appendix B.3).

B.2 Correctness

Let co_k^{ij} be the amount agent $i \in [m]$ is assigned for consumption from $j \in [m]$ of $k \in [\ell]$ at the old price ($\frac{p_k^j}{(1+\epsilon)}$) and cn_k^{ij} be the amount agent i is assigned for consumption from j of k at the new price p_k^j . Similarly, let ro_k^{ij} be the amount agent i is assigned for resale from j of k at the old price ($\frac{p_k^j}{(1+\epsilon)}$) and rn_k^{ij} be the amount agent i is assigned for resale from j of k at the new price p_k^j . Then $x_k^{ij} = co_k^{ij} + cn_k^{ij}$ and $y_k^{ij} = ro_k^{ij} + rn_k^{ij}$. Define the surplus budget with agent i as

$$s_i = \sum_{k \in [\ell]} p_k^i e_k^i + \sum_{j=i} \sum_{k \in [\ell]} p_k^i (ro_k^{ij} + rn_k^{ij}) - \sum_{j=i} \sum_{k \in [\ell]} p_k^j \left(\frac{ro_k^{ij}}{1+\epsilon} + rn_k^{ij} \right) - \sum_{j=i} \sum_{k \in [\ell]} p_k^j$$

and the total surplus budget in the economy as $s = \sum_{i \in [m]} s_i$. Finally, let $e_{min} = \min_{i,k} \{e_k^i : e_k^i > 0\}$, $b_{min} = \min_i \{b^i : b^i > 0\}$, and $\zeta = \min(e_{min}, \frac{\epsilon}{1+\epsilon} b_{min})$. Note that in the definition of s_i , agent i is never reselling good k at their old price for the good

ALGORITHM 1: Auction Based Approximation Algorithm (Part 1)

algorithm Main:

```
initialization:  $\forall i, j \in [m]: \mathbf{x}^i = \mathbf{y}^i = \mathbf{0}, \mathbf{p}^i = \mathbf{1},$   
                   $\beta_i = s_i = \mathbf{p}^i \cdot \mathbf{e}^i$   
while  $\sum_{i \in [m]} s_i > \frac{\epsilon}{1+\epsilon} \zeta$  or  $\exists i \in [m], k \in [\ell]: \sum_{j \in [m]} x_k^{ij} + y_k^{ij} \neq$   
           $e_k^i + \sum_{j \in [m]} y_k^{ij}$  do  
    pick  $i \in [m]$  such that  $s_i > 0$   
    let  $\hat{\mathbf{x}}^i = \mathcal{O}_{C_i}(\mathbf{p}, \beta_i, \mathbf{x}^i)$   
    update  $\mathbf{r} = \hat{\mathbf{x}}^i - \mathbf{x}^i$  // compute consumption request  
    vectors  
    if Assign( $i, \mathbf{x}^i, \mathbf{r}$ ) // neighbors have unassigned  
    requested good(s)  
    else if Outbid( $i, \mathbf{x}^i, \mathbf{r}$ ) // neighbors assigned  
    requested good(s) at lower price  
    else Reschedule_Resale( $i, \mathbf{x}^i, \mathbf{r}$ ) // neighbors try to  
    get more requested good(s)  
end
```

procedure Assign($i, \mathbf{z}, \mathbf{r}$):

```
let  $\mathbf{a}^{total} = \mathbf{0}$  // tracks good(s) assigned, vector of  
size  $\ell$   
for  $j \approx i$  and  $k \in [\ell]$  such that  $r_k^j > 0$  do  
  let  $a$  be the minimum of: (i) the amount  $j$  has unassigned of  $k$   
  and (ii)  $r_k^j$   
  update  $z_k^j = z_k^j + a$  at price  $p_k^j$   
  if  $\mathbf{z}$  is a consumption plan then update  $s_i = s_i - ap_k^j$   
  update  $a_k^{total} = a_k^{total} + a$   
end  
return  $\mathbf{a}^{total}$ 
```

procedure Outbid($i, \mathbf{z}, \mathbf{r}$):

```
let  $\mathbf{a}^{total} = \mathbf{0}$  // tracks good(s) assigned, vector of  
size  $\ell$   
while  $r_k^j > 0$  and  $j \approx i$  has neighbor  $\hat{j} \approx j$  with  $k \in [\ell]$  assigned at  
a lower price do  
  let  $c$  be the amount  $\hat{j}$  bought of  $k$  at a lower price for  
  consumption  
  let  $d$  be the amount  $\hat{j}$  bought of  $k$  at a lower price for resale  
  let  $a = \min(c + d, r_k^j)$   
  update  $z_k^j = z_k^j + a$  at price  $p_k^j$   
  if  $\mathbf{z}$  is a consumption plan then update  $s_i = s_i - ap_k^j$   
  unassign  $a$  units of  $k$  from  $\hat{j}$ , from consumption and then  
  resale  
  update  $s_j = s_j + \min(a, c)p_k^j/(1 + \epsilon)$   
  if  $a - c > 0$  then Update_Resale( $\hat{j}, k, a - c$ )  
  update  $a_k^{total} = a_k^{total} + a$   
end  
return  $\mathbf{a}^{total}$ 
```

ALGORITHM 1: Auction Based Approximation Algorithm (Part 2)

procedure Reschedule_Resale($i, \mathbf{z}, \mathbf{r}$):

```
let  $\mathbf{a}^{total} = \mathbf{0}$  // tracks good(s) assigned, vector of  
size  $\ell$   
for  $j \approx i$  such that  $\exists r_k^j > 0$  for some  $k \in [\ell]$  do  
  let  $\hat{\mathbf{y}}^j = \mathcal{O}_{R_j}(\mathbf{p}, b_j, \mathbf{y}^j, \mathbf{r}^j)$   
  update  $\hat{\mathbf{r}} = \hat{\mathbf{y}}^j - \mathbf{y}^j$  // compute resale request vectors  
  let  $\mathbf{a} = \min(\sum_{j \approx i} \hat{\mathbf{r}}^j, \mathbf{r}^j)$  // min is taken elementwise  
  initialize  $\mathbf{q} = \mathbf{0}$  // book keeping vector  
  for Function  $\in \{\text{Assign, Outbid, Reschedule\_Resale}\}$  and  
  if  $\mathbf{q} < \mathbf{a}$  do // In order  
     $\hat{\mathbf{q}} = \text{Function}(j, \mathbf{y}^j, \mathbf{a})$  // Make assignment, track  
    amount  
    update  $\mathbf{q} = \mathbf{q} + \hat{\mathbf{q}}$   
  end  
  for  $\hat{j} \approx j$  and  $k \in [\ell]$  such that  $\hat{r}_k^{\hat{j}} < 0$  do  
    let  $c$  and  $d$  be the amounts  $\hat{j}$  bought of  $k$  for consumption  
    and resale respectively  
    unassign  $\hat{r}_k^{\hat{j}}$  units of  $k$  from  $\hat{j}$ , from consumption and  
    then resale  
    if  $\hat{r}_k^{\hat{j}} - c > 0$  then Update_Resale( $\hat{j}, k, \hat{r}_k^{\hat{j}} - c$ )  
  end  
  update  $\mathbf{z}^j = \mathbf{z}^j + \mathbf{q}$  at prices  $\mathbf{p}^j$   
  if  $\mathbf{z}$  is a consumption plan then update  $s_i = s_i - \mathbf{q} \cdot \mathbf{p}^j$   
  update  $\mathbf{a}^{total} = \mathbf{a}^{total} + \mathbf{q}$   
  for  $k \in [\ell]$  if  $\exists a_k - q_k > 0$  or  $a_k < r_k^j$  do  
    Raise_Price( $j, k$ )  
end  
return  $\mathbf{a}^{total}$ 
```

procedure Update_Resale(j, k, q):

```
for  $i \in [m]$  such that  $x_k^{ij} > 0$  and while  $q > 0$  do  
  let  $a = \min(q, x_k^{ij})$   
  update  $x_k^{ij} = x_k^{ij} - a$   
  update  $s_i = s_i + ap_k^j/(1 + \epsilon)$   
  update  $q = q - a$   
end  
for  $i \in [m]$  such that  $y_k^{ij} > 0$  and while  $q > 0$  do  
  let  $a = \min(q, y_k^{ij})$   
  update  $y_k^{ij} = y_k^{ij} - a$   
  Update_Resale( $i, k, a$ )  
  update  $q = q - a$   
end
```

procedure Raise_Price(j, k):

```
update  $s_j = s_j + \epsilon p_k^j e_k^j$   
 $p_k^j = p_k^j (1 + \epsilon)$   
for  $\hat{j} \approx j$  if  $b_j > 0$  do  
  let  $\hat{\mathbf{y}}^{\hat{j}} = \mathcal{O}_{R_{\hat{j}}}(\mathbf{p}, b_{\hat{j}}, \mathbf{y}^{\hat{j}}, \mathbf{1})$  // consult oracle for  
  optimal profits, arbitrary request  
  update  $s_j$  so that resale profits are  $\sum_{\hat{j} \approx j} (\mathbf{p}^{\hat{j}} - \mathbf{p}^j) \cdot \hat{\mathbf{y}}^{\hat{j}}$   
  while  $s_j < 0$  do unassign goods from  $\mathbf{x}^{\hat{j}}$   
end
```

$\frac{p_k^i}{(1+\epsilon)}$. This follows from the fact that *Raise_Rrice* can only be called after an agent has rescheduled as much resale as they can.

We will show in Lemma B.1 that the following hold throughout a run of the algorithm. As in Definition 5.2, here $\bar{\mathbf{p}}$ satisfies $\bar{\mathbf{p}}^i = \mathbf{p}^i$ and $\bar{\mathbf{p}}^j = \mathbf{p}^j / (1 + \epsilon)$ for every $j \in [m] \setminus \{i\}$, and we define budgets $\beta_i = \mathbf{p}^i \cdot \mathbf{e}^i + \sum_{j \neq i} (\mathbf{p}^i - \mathbf{p}^j) \cdot \mathbf{y}^{ij}$ and $\bar{\beta}_i = \bar{\mathbf{p}}^i \cdot \mathbf{e}^i + \sum_{j \neq i} (\bar{\mathbf{p}}^i - \bar{\mathbf{p}}^j) \cdot \mathbf{y}^{ij}$.

Invariant 1. $\mathbf{y}^i \leq \bar{\mathbf{y}}^i$ for some $\bar{\mathbf{y}}^i \in R_i(\bar{\mathbf{p}}, b_i)$.

Invariant 2. $\mathbf{x}^i \leq \bar{\mathbf{x}}^i$ for some $\bar{\mathbf{x}}^i \in C_i\left(\frac{\mathbf{p}}{1+\epsilon}, \bar{\beta}_i\right)$.

Invariant 3. $\sum_{j \neq i} x_k^{ji} + \sum_{j \neq i} y_k^{ji} \leq e_k^i + \sum_{j \neq i} y_k^{ij}$ for all $k \in [\ell]$.

Invariant 4. If $e_k^i = 0$, then $\sum_{j \neq i} y_k^{ij} = \sum_{j \neq i} (y_k^{ji} + x_k^{ji})$.

Invariant 5. $\sum_{j \neq i} \mathbf{p}^j \cdot \mathbf{x}^{ij} \leq \beta_i$ for all $i \in [m]$.

LEMMA B.1. *At every iteration of the main loop in Algorithm 1, for every agent $i \in [m]$, neighbor $j \neq i$, and good $k \in [\ell]$, Invariants 1 through 5 hold.*

PROOF. To derive Invariant 1, notice that agents are only ever reselling goods at prices $\hat{\mathbf{p}}$ satisfying $\bar{\mathbf{p}} \leq \hat{\mathbf{p}} \leq \mathbf{p}$ after consulting their resale oracle. The agent $i \in [m]$ has a constant resale budget b_i . Recall that resale demand systems satisfy WGS for $-R_i$ and so, by the definition of WGS, for $-\hat{\mathbf{y}}^i \in -R_i(\hat{\mathbf{p}}, b_i)$ and $-\bar{\mathbf{y}}^i \in -R_i(\bar{\mathbf{p}}, b_i)$ we know that $-\hat{y}_k^{ij} \geq -\bar{y}_k^{ij}$ whenever $\bar{p}_k^j \leq \hat{p}_k^j$ for all $j \in [m], k \in [\ell]$. Thus $\hat{y}_k^{ij} \leq \bar{y}_k^{ij}$ for all j and k , which implies that $\hat{\mathbf{y}}^i \leq \bar{\mathbf{y}}^i$. We know that $\hat{\mathbf{y}}^i$ is the resale plan obtained during i 's last call to their resale oracle, and throughout the algorithm goods might be unassigned, so it follows that at any point in the algorithm $\mathbf{y}^i \leq \hat{\mathbf{y}}^i \leq \bar{\mathbf{y}}^i$ and Invariant 1 holds.

Invariant 2 holds for a similar reason. Notice that agents are only ever consuming and reselling goods at prices $\hat{\mathbf{p}}$ satisfying $\frac{\mathbf{p}}{1+\epsilon} \leq \bar{\mathbf{p}} \leq \hat{\mathbf{p}} \leq \mathbf{p}$ after consulting their demand oracle. Let $i \in [m]$ be an agent, $\bar{\beta}_i = \bar{\mathbf{p}}^i \cdot \mathbf{e}^i + \sum_{j \neq i} (\bar{\mathbf{p}}^i - \bar{\mathbf{p}}^j) \cdot \mathbf{y}^{ij}$ and $\hat{\beta}_i = \hat{\mathbf{p}}^i \cdot \mathbf{e}^i + \sum_{j \neq i} (\hat{\mathbf{p}}^i - \hat{\mathbf{p}}^j) \cdot \mathbf{y}^{ij}$. As agents never resell their own goods at an old price by design, i 's budget at the price set $\bar{\mathbf{p}}$ is maximal at the iteration of the algorithm being considered, i.e. $\bar{\beta}_i \geq \hat{\beta}_i$. By the definition of WGS, for $\bar{\mathbf{x}}^i \in C_i\left(\frac{\mathbf{p}}{1+\epsilon}, \bar{\beta}_i\right)$ and $\hat{\mathbf{x}}^i \in C_i(\mathbf{p}, \hat{\beta}_i)$ we know that $\hat{x}_k^{ij} \leq \bar{x}_k^{ij}$ whenever $\bar{p}_k^j \leq \hat{p}_k^j$ for all $j \in [m], k \in [\ell]$. Thus we know $\hat{\mathbf{x}}^i \leq \bar{\mathbf{x}}^i$. As $\hat{\mathbf{x}}^i$ is the consumption plan obtained during i 's last call to their demand oracle, and goods might be unassigned during the algorithm, we know that $\mathbf{x}^i \leq \hat{\mathbf{x}}^i \leq \bar{\mathbf{x}}^i$ and Invariant 2 holds.

Invariant 3 is true because none of the procedures in the algorithm ever over-allocate goods and we carefully propagate changes to resale plans throughout the economy when they appear.

Invariant 4 holds because in Algorithm 1 an agent only ever has a good assigned for resale if an agent is bidding on it from them and they are unable to provide it from endowments. If an agent no longer wants the good being resold to them, then that fact is propagated throughout the economy in *Update_Resale* and the agent returns the good to their supplier. Thus if the agent reselling a good is endowed with it then the inequality is strict, and if the agent is not endowed with the good then equality holds.

Invariant 5 is true simply because agents never place bids that bring their surplus budget below zero and when a surplus budget falls below zero (due to the Outbid procedure taking back goods being resold) Algorithm 1 ensures that goods get unassigned until the surplus becomes zero again. \square

From Invariants 1 and 2 we are guaranteed to have approximately optimal arbitrage and approximate individual rationality satisfied at termination. Invariants 3 and 5 partially imply approximate clearing and approximately efficient budget use at termination. In fact, combining Invariants 3 and 4 shows us that when $\mathbf{e}^i = \mathbf{0}$ local clearing condition is exact (i.e. not approximate) at termination. It remains to show that approximate clearing and approximately efficient budget use are bounded from below, which we establish in the following two lemmas.

LEMMA B.2. *When Algorithm 1 terminates:*

$$\forall i \in [m], k \in [\ell]: \frac{e_k^i + \sum_{j \neq i} y_k^{ij}}{(1+\epsilon)} \leq \sum_{j \neq i} x_k^{ji} + \sum_{j \neq i} y_k^{ji} \leq e_k^i + \sum_{j \neq i} y_k^{ij} \quad (5)$$

PROOF. If the algorithm terminates because all goods clear locally in the economy, then the right inequality is an equality, and the left inequality follows trivially. Now consider the other termination case, when the total surplus budget in the economy is at most $\frac{\epsilon}{1+\epsilon} e_{min}$. In light of Invariant 4, we need only prove the left inequality. Using the fact that $1 - \frac{1}{1+\epsilon} = \frac{\epsilon}{1+\epsilon}$ allows us to rewrite the surplus budget of agent i as

$$s_i = \sum_{k \in [\ell]} p_k^i e_k^i + \sum_{j \neq i} \sum_{k \in [\ell]} p_k^j y_k^{ij} - \sum_{j \neq i} \sum_{k \in [\ell]} p_k^j y_k^{ji} - \sum_{j \neq i} \sum_{k \in [\ell]} p_k^j x_k^{ji} + \frac{\epsilon}{1+\epsilon} \sum_{j \neq i} \sum_{k \in [\ell]} p_k^j$$

Combining the termination condition with the observations $p_k^i \geq 1$, $y_k^{ij} \geq 0$, $x_k^{ij} \geq 0$, $ro_k^{ij} \geq 0$, and $co_k^{ij} \geq 0$, we then have

$$\sum_{i \in [m], k \in [\ell]} \left[e_k^i + \sum_{j \neq i} (y_k^{ij} - y_k^{ji} - x_k^{ji}) \right] \leq \sum_{i \in [m]} s_i \leq \frac{\epsilon}{1+\epsilon} e_{min}.$$

As throughout the algorithm, goods are never allocated more than they can be supplied, it follows that all summands of the outer sum are non-negative, implying that

$$\forall i \in [m], k \in [\ell]: e_k^i + \sum_{j \neq i} y_k^{ij} - \sum_{j \neq i} (y_k^{ji} + x_k^{ji}) \leq \frac{\epsilon}{1+\epsilon} e_{min}.$$

Now fix i and k and consider two cases; in the first, suppose $e_k^i > 0$. Then by definition of e_{min} , we have $e_{min} \leq e_k^i + \sum_{j \neq i} y_k^{ij}$ and thus after rearranging terms, $\frac{e_k^i + \sum_{j \neq i} y_k^{ij}}{(1+\epsilon)} \leq \sum_{j \neq i} x_k^{ji} + \sum_{j \neq i} y_k^{ji}$. In the second case, suppose $e_k^i = 0$. Then by Invariant 5 we have $\sum_{j \neq i} y_k^{ij} = \sum_{j \neq i} (y_k^{ji} + x_k^{ji})$ which immediately implies the bound. \square

LEMMA B.3. *When Algorithm 1 terminates:*

$$\forall i \in [m]: \frac{\mathbf{p}^i \cdot \mathbf{e}^i + \sum_{j \neq i} (\mathbf{p}^i - \mathbf{p}^j) \cdot \mathbf{y}^{ij}}{1+\epsilon} \leq \sum_{j \neq i} \mathbf{p}^j \cdot \mathbf{x}^{ij} \leq \mathbf{p}^i \cdot \mathbf{e}^i + \sum_{j \neq i} (\mathbf{p}^i - \mathbf{p}^j) \cdot \mathbf{y}^{ij}$$

PROOF. The right inequality holds by Invariant 5. Furthermore, notice that if the algorithm terminates because all goods clear locally in the economy, then this implies that no good is being over demanded which can only happen if all agents are able to spend their budgets entirely. Thus the left inequality follows trivially. Consider the other termination case, when the total surplus in the economy $s \leq \frac{\epsilon}{1+\epsilon} \zeta$. From the definition of agent i 's surplus we get that

$$s_i = \sum_{k \in [\ell]} p_k^i e_k^i + \sum_{j \neq i} \sum_{k \in [\ell]} p_k^i y_k^{ij} - \sum_{j \neq i} \sum_{k \in [\ell]} p_k^j y_k^{ij} - \sum_{j \neq i} \sum_{k \in [\ell]} p_k^j x_k^{ij} + \frac{\epsilon}{1+\epsilon} \sum_{j \neq i} \sum_{k \in [\ell]} p_k^j (e_k^{ij} + M(R_i))$$

which implies that

$$p^i \cdot e^i + \sum_{j \neq i} (p^i - p^j) \cdot y^{ij} - \frac{\epsilon}{1+\epsilon} \zeta \leq \sum_{j \neq i} p^j \cdot x^{ij}.$$

Now fix i and consider two cases; in the first, suppose there exists $k \in [\ell]$ such that $e_k^i > 0$. By definition we know $p^i \cdot e^i \geq \zeta$ (since $p_k^i \geq 1$) and thus after rearranging terms, $\frac{p^i \cdot e^i + \sum_{j \neq i} (p^i - p^j) \cdot y^{ij}}{1+\epsilon} \leq \sum_{j \neq i} p^j \cdot x^{ij}$. In the second case, suppose $e_k^i = 0$ for every $k \in [\ell]$. Note that for any $j \in [m]$ and $k \in [\ell]$ such that $y_k^{ij} > 0$, we know $p_k^i - p_k^j \geq \frac{\epsilon}{1+\epsilon}$ and so $\sum_{j \neq i} (p^i - p^j) \cdot y^{ij} \geq \zeta$. Again, rearranging terms implies that $\frac{p^i \cdot e^i + \sum_{j \neq i} (p^i - p^j) \cdot y^{ij}}{1+\epsilon} \leq \sum_{j \neq i} p^j \cdot x^{ij}$ which concludes the proof. \square

Thus by Invariants 1 and 2 along with Lemmas B.2 and B.3, we know that when Algorithm 1 terminates it has found a $(1 + \epsilon)$ -approximate b -resale equilibrium.

B.3 Complexity Analysis

Let p_{max} be the largest price found by Algorithm 1. As highlighted in §5.3, we will use p_{max} to: (i) bound the number of calls made to the *Raise_Price* procedure during the algorithm (Lemma B.4) and (ii) bound the amount of time between calls to *Raise_Price* before the algorithm terminates (Lemma B.5). An upper bound on p_{max} can be found for specific sets of demand and resale systems, i.e. when these systems are given by explicit utility functions and resale constraints.

LEMMA B.4. *The number of calls to the *Raise_Price* procedure is bounded by:*

$$O\left(\frac{\ell}{\epsilon} \log(p_{max})\right)$$

PROOF. By design, prices are never decreased in Algorithm 1. Furthermore, each call to the *Raise_Price* procedure raises the price of a single good by a factor of $(1 + \epsilon)$. Thus the number of price raises is bounded by $\ell \log_{1+\epsilon} p_{max}$. \square

We further partition the steps of the algorithm into *rounds*. A single round corresponds to a sequence of iterations of the while loop in Algorithm 1's Main function, such that each agent i with surplus budget $s_i > 0$ has reduced s_i to 0 at least once after calling the Assign, Outbid, or Reschedule_Resale procedures (possibly interleaved with other agents' calls to these functions). We say a round is *complete* when the *Raise_Price* procedure is not called during the round. It is important to note that, due to the Outbid and

Reschedule_Resale procedures, by the end of a round an agent's surplus s_i may become positive again.

Recall that $M(R_i)$ was defined to be the maximal quantity of resold goods $\|y^i\|_1$ for any $y^i \in R_i(p^*, b_i)$ and any p^* is $(0, p_{max}]^{m \times \ell}$.

LEMMA B.5. *In a sequence of at least $\frac{(1+\epsilon)(p_{max} \sum_{i \in [m], k \in [\ell]} (e_k^i + M(R_i)))}{\epsilon e_{min}}$ complete rounds, either a call to *Raise_Price* is made or Algorithm 1 terminates.*

PROOF. Consider a sequence of complete rounds in which no calls to *Raise_Price* are made, noting that all prices in the economy are the same during such a sequence. Let the first round in the sequence be referred to as round t_0 , the second round in the sequence as round $t_0 + 1$, and so on. Throughout this proof we will adopt the convention that $t \geq t_0$ and round t is still part of the sequence of complete rounds being considered. Let $s(t)$ be the total surplus budget at the beginning of round t . Define the total amount of money spent on goods at their old price $\frac{p_k^j}{1+\epsilon}$ throughout the economy at the start of round t as $\tau(t) = \sum_{i, j \in [m], k \in [\ell]} \frac{p_k^j}{1+\epsilon} (co_k^{ij}(t) + ro_k^{ij}(t))$ where $co_k^{ij}(t)$ (resp., $ro_k^{ij}(t)$) is the amount agent i spent on agent j 's good k at the old price for consumption (resp., resale) at the beginning of round t .

Observe that for the sequence to continue into another round, the calls to Outbid and Reschedule_Resale must return at least $\frac{\epsilon}{1+\epsilon} e_{min}$ amount of money to the total surplus in the economy by the end of the round or else the algorithm terminates. Let $\mu(t)$ be the total amount of money returned to surplus budgets via calls to Outbid and Reschedule_Resale during round t and let $\mu_0(t)$ be the contributions to $\mu(t)$ stemming from increases to s_i after an agent i had $s_i = 0$ during round t . As every agent achieves zero surplus budget at some point during the round, by definition, we in fact have $\mu_0(t) = s(t+1)$. We therefore must have $\mu(t) \geq \mu_0(t) = s(t) > \frac{\epsilon}{1+\epsilon} e_{min}$, as the algorithm has not terminated. In addition, Algorithm 1 never assigns goods at the old prices and can only unassign goods if they were previously assigned at the old price (i.e. neither Outbid nor Reschedule_Resale can unassign a good bought at the current price p_k^j). This implies that all money in $\mu(t)$ was previously spent on goods at old price and all money in $\mu(t)$ will be spent to outbid goods bought at old prices. Therefore $\tau(t_0) > \tau(t)$ and we must have $\tau(t+1) = \tau(t) - \mu(t)$. Once we reach round t in the sequence where $\tau(t) = 0$, we know that in the round $t+1$ either a call to *Raise_Price* is made or the algorithm will terminate. This follows from the fact that, when Reschedule_Resale is called, it will call *Raise_Price* if and only if it is not able to supply the requested amount of goods. Since $\tau(t) = 0$, in round $t+1$ either every call to Reschedule_Resale is able to meet demand (satisfying the local clearing condition for termination) or a good is over demanded and it calls *Raise_Price*. Thus, we must have $t - t_0 + 1 \leq \frac{\tau(t_0)}{\frac{\epsilon}{1+\epsilon} e_{min}}$.

As the total consumption is bounded by the sum of endowments, the total resale by the sum of $M(R_i)$, and prices by p_{max} , we must have $\tau(t_0) \leq p_{max} \sum_{i \in [m], k \in [\ell]} (e_k^i + M(R_i))$. Therefore $t - t_0 + 1 \leq \frac{(1+\epsilon)(p_{max} \sum_{i \in [m], k \in [\ell]} (e_k^i + M(R_i)))}{\epsilon e_{min}}$. \square

Having shown that Algorithm 1 finds a $(1 + \epsilon)$ -approximate b -resale equilibrium at termination in Appendix B.2, we now conclude

the proof of Theorem 5.3 by formally deriving the time complexity of Algorithm 1. In Lemma B.4 we bounded the maximum number of calls to the *Raise Price* procedure during the algorithm, and in Lemma B.5 we bounded the number of rounds between calls to *Raise Price*, so all that remains is to derive the time complexity of a single round. That said, the time spent in a round is fully determined by the number of calls to the Assign, Outbid, and Reschedule_Resale procedures during the round. However, the pseudo-code presented for Algorithm 1 above is purposely written in a way that highlights its decentralized and asynchronous nature. Therefore, for the sake of deriving a time complexity for Theorem 5.3, we will assume that in a round: (i) agents are considered in an arbitrary but fixed order and (ii) while that agent’s surplus budget is greater than zero the agent will repeat the process described in *Main*’s loop. Note that if the agent chosen in (ii) already has no surplus budget, this will count as “reducing s_i to zero at least once” in that round.

Consider an agent i during a complete round that proceeds in the manner just described. Each time i calls the Assign procedure they are asking neighbors to give them as much of each desired good as the neighbor currently have unassigned. Each time i calls the Outbid procedure they are asking neighbors to re-neg on deals they have already made on the desired goods. Each time i calls the Reschedule_Resale procedure they are asking neighbors to do whatever they can to resell desired goods to them. Thus, since Algorithm 1 never unassigns goods that are currently assigned at the current price, the time needed for each of these procedures (without forcing a call to *Raise Price*) is bounded by exactly the time needed for each procedure to check every source of each good (and to query resale oracles in the case of Reschedule_Resale). To formally derive this bound, let $N_i = |\{j \simeq i\}|$ be the size of i ’s open neighborhood in the economy’s underlying graph and let $N_{max} = \max_{i \in [m]} N_i$.


During the start of its turn, agent i will make a single call to its consumption oracle which takes time T_D . Agent i then, in the worst case, makes a call to Assign, Outbid, and also Reschedule_Resale. Each call to Assign essentially amounts to book-keeping and updating variables. The agent checks each $j \simeq i$ per good $k \in [\ell]$, for a total of $N_i \ell$ operations in the worst case. Thus the time needed for Assign is $O(N_{max} \ell)$. Calls to Outbid require some book-keeping along with any time spent in the *Update_Resale* procedure. In the worst case, during this procedure the agent checks each $\hat{j} \simeq j$ per $j \simeq i$ per good $k \in [\ell]$ and can make one call to *Update_Resale* each time. Notice that the *Update_Resale* procedure is essentially a breadth first search with some variables being updated along the way, so we know its time complexity to be $O(mN_{max})$. Thus the Outbid procedure needs time at most $O(N_{max}^3 m \ell)$. Calls to Reschedule_Resale requires calls to resale oracles, some book-keeping, internal calls to Assign and Outbid, and also recursive calls to Reschedule_Resale. We know calls to resale oracles take time T_R and each neighbor $j \simeq i$ must consult their oracle in the worst case. The time complexity of each internal call to Assign and Outbid is the same as we’ve already derived above, therefore all we have left to derive is the number of recursive calls to Reschedule_Resale that are possible. Note that Reschedule_Resale is only called recursively for neighbors whose goods are profitable to resell, therefore we know that cycles are impossible. It follows that the depth of these recursive

calls to Reschedule_Resale is bounded by m . Within each call to Reschedule_Resale internal calls to Assign and Outbid can again be made. From this we know that the Reschedule_Resale procedure has time complexity $O(N_{max} m T_R + N_{max}^3 m^2 \ell)$.

Combining the above, the time spent on a single agent is bounded by $O(T_D + N_{max} m T_R + N_{max}^3 m^2 \ell)$ and so the time complexity of a single round is bounded by $O(m T_D + N_{max} m^2 T_R + N_{max}^3 m^3 \ell)$. Combining this bound with the lemmas found in this section proves Theorem 5.3.

C ASYMMETRIC BROKER EXAMPLE

We will demonstrate that addressing non-existence of equilibria by imposing a small minimum on endowments (i.e. replacing each zero in endowment vectors with ϵ) is not a suitable alternative to resale in graphical economies. To do so, consider the following modified 3-node broker example highlighted in §3. Again, for every agent $i \in [3]$, the consumption demand systems C_i given by linear utilities u^i and R_i is a credit bound resale demand system.




We show that no KKO equilibrium exists (§C.1), and that any KKO equilibrium arising from imposing a small minimum on endowments—the ϵ -KKO equilibria—wastes an entire unit of utility (§C.2). By contrast, we conclude by showing that incorporating resale gives rise to an equilibrium matching economic intuition and settles this tension borne from KKO being too local (§C.3).

C.1 Non-existence of KKO Equilibria

As before, two agents with complementary endowments and preferences are connected only to a single broker agent with no endowment; the only difference is that the broker agent has a singular preference for one of the goods. Given the similarity to the original broker example, an essentially identical argument to the one given in §3 reveals that no KKO equilibrium exists. To summarize: Without resale the market cannot clear if any agent purchases goods from agent 2 as $e^2 = (0, 0)$. Therefore $p_1^1 = p_2^3 = 0$, otherwise agents 1 or 3 have non-zero budget and by individual rationality must spend it on goods their neighbors are not endowed with. However, now that these prices are zero, agent 2 only satisfies individual rationality by buying an infinite amount of the second good. Thus no KKO equilibrium exists.

C.2 Wasteful Outcomes at ϵ -KKO Equilibria

Suppose a small minimum is imposed on agents’ endowments, and all zeros are replaced with ϵ in the endowments. We argue that the resultant ϵ -KKO equilibria must have the following prices up to uniformly re-scaling prices by a multiplicative constant $\gamma > 0$. Dotted lines depict one possible movement of goods, but we show all other outcomes are equivalent.



Since we are interested in ϵ -KKO equilibrium outcomes, no resale is possible and goods can only be traded with immediate neighbors. Therefore agents 1 and 2 are the only possible consumers for agent 1's endowment of the first good e_1^1 , but neither agents 1 or 2 have utility for this good. All endowments must be consumed for the market to clear in eq. (3), so it follows that $p_1^1 = 0$ or else by individual rationality agents 1 and 2 will not consume the good. Furthermore, we know $p_2^1 = p_2^2 = p_3^3$ otherwise market clearing is impossible since individual rationality dictates that goods are only ever consumed at the lowest price in an agent's neighborhood. For the same reason we get that $p_2^1 = p_3^1$.


Let $\alpha = p_2^1 = p_2^2 = p_3^3$ and $\beta = p_1^2 = p_1^3$, where we know $\alpha, \beta > 0$ otherwise individual rationality is only satisfied by agents consuming an infinite amount of the goods they have utility for. We find that agents 1, 2, and 3 sell their endowments for a revenue of $p^1 \cdot e^1 = \epsilon\alpha$, $p^2 \cdot e^2 = \epsilon\beta + \epsilon\alpha$, and $p^3 \cdot e^3 = \beta + \epsilon\alpha$ respectively. For simplicity, let us assume that agent 1 spends its entire $\epsilon\alpha$ revenue to purchase back its own endowment vector for consumption.⁷ In this case, agent 2 spends its entire revenue to purchase the entirety of both its own and agent 3's endowment of the second good, e_2^2 and e_2^3 respectively—it is intuitive to think of agent 2 using the revenue $\epsilon\beta$ from selling e_1^2 to purchase e_2^3 and the revenue $\epsilon\alpha$ from selling e_2^2 to purchase itself back. Clearly then agent 3 spends its entire revenue to purchase the entirety of both its own and agent 2's endowment of the first good, e_1^3 and e_1^2 respectively. At the end of this process, agents 1, 2, and 3 have consumption plans summarized by the bundles $(1, \epsilon)$, $(0, 1 + \epsilon)$, and $(2\epsilon, 0)$ respectively.

Local market clearing and individual rationality are clear since we ensured all endowments are consumed by immediate neighbors, agents only spend revenue on utility maximizing goods, and each agent spends exactly their entire revenue. To derive the prices shown above, let $\beta = 1$. From agent 2's budget constraint we find that $\epsilon\alpha + \alpha = \epsilon + \epsilon\alpha$, where the left hand side is the cost of agent 2's consumption bundle and the right hand side is their revenue. Hence, when $\beta = p_1^2 = p_1^3 = 1$ we get $\alpha = p_2^1 = p_2^2 = p_3^3 = \epsilon$. By rescaling β , it is clear that all other ϵ -KKO equilibrium prices must be a rescaling of these prices by a multiplicative constant. In addition, we saw that an entire unit of utility is lost at ϵ -KKO equilibria and the final utilities are given by ϵ , $1 + \epsilon$, and 2ϵ for agents 1, 2, and 3 respectively.

⁷We have equivalent outcomes when agents 1 and 2 "split" e_1^1 and/or "swap" consumption from equal parts of e_1^1 and e_2^2 . The reason e_1^1 can be "split" at equilibrium is because it is free and neither agent has utility for the good. Similarly, e_1^2 and e_2^2 can have equal parts "swapped" between the agents because the good is equally priced by both agents and they have identical utility for the good.

C.3 Efficient Outcomes at Resale Equilibria

We now conclude by showing that allowing resale leads to equilibria that match economic intuition even when endowment vectors are sparse. We will argue that, for any $0 < b \leq 2$, the following prices at $\alpha = \frac{\sqrt{1+4b}-1}{2}$ lead to a b -resale equilibrium where $b = (b, b, b)$. Observe that $\alpha = \frac{\sqrt{1+4b}-1}{2}$ is a solution of the quadratic equation for the polynomial $\alpha^2 + \alpha = b$.



To verify the equilibrium, recall that it is helpful to think of two phases. Agents 1 and 3 have nothing to gain from resale, so abstain from the first phase, whereas agent 2 can profit and is willing to resell any bundle of goods with a total cost of b so long as the goods being resold maximize profit-per-credit. By definition $0 < \alpha \leq 1$ and $p_1^2/p_1^1 = p_2^2/p_2^3 = 1/\alpha$, therefore the goods with non-zero endowments always have maximal profit-per-credit and agent 2 can resell any combination of these goods with a total cost of b . Furthermore, as all prices are non-zero and individual rationality dictates that agents will not spend any budget on goods they have no utility for, we know that agent 2 must purchase $(1, 0)$ from agent 1 for resale in order for the market to clear. It follows that, since $p_1^1 = \alpha$, $p_2^3 = 1$, and the total cost of goods resold must be b , agent 2 purchases $(0, b - \alpha) = (0, \alpha^2)$ from agent 3 for resale. Hence, agent 2 makes an optimal profit from resale of $1 - \alpha^2$ and agent 1 has sold its entire endowment for a total revenue of α . In the second phase, agent 3 sells its remaining endowment, which combines with the revenue made during the resale phase, for a total revenue of 1. Agents 1 and 2 are only interested in buying the second good and will spend their entire budget doing so at the best bang-per-buck price available; agent 1 thus purchases $(0, \alpha^2)$ from agent 2 at a cost of α and agent 2 purchases $(0, 1 - \alpha^2)$ from agent 3 at a cost of $1 - \alpha^2$. Similarly, agent 3 is only interested in buying the first good and will spend their entire budget doing so; thus agent 3 purchases $(1, 0)$ from agent 2 at a cost of 1.

Optimal arbitrage and individual rationality are clear since agents only resold maximal profit-per-credit goods and bought maximal bang-per-buck goods using their entire budgets; it remains to check the clearing constraint. We know agent 2 has an endowment of $(0, 0)$ and resold $(1, \alpha^2)$, therefore since agents 1 and 3 consumed $(0, \alpha^2)$ and $(1, 0)$ from agent 2, respectively, the market locally clears at agent 2. Similarly, agents 1 and 3 had endowments of $(1, 0)$ and $(0, 1)$ respectively. Since agent 2 resold $(1, 0)$ and $(0, \alpha^2)$ from agents 1 and 3, respectively, and consumed $(0, 1 - \alpha^2)$ from agent 3, it is clear that the market locally clears at agents 1 and 3. Therefore we have confirmed that this is a b -resale equilibrium.

Thus adding resale resolves this modified broker example in an intuitive way: adding a small amount of resale capacity b facilitates trade between agents 1 and 3, while allowing agent 2 to extract rent

and consume nearly all of the good it has utility for. In particular, the final utilities are α^2 , $1 - \alpha^2$, and 1, respectively. The larger b becomes, the less rent agent 2 can extract, for the simple reason that the prices for the other agents must increase for the market to clear. When $b \geq 2$, all prices become equal and the market is effectively a classic AD exchange economy as predicted in §3.1. Regardless of b , we see that the social welfare is strictly better than the ϵ -KKO equilibrium outcomes found in §C.2 since all utility is extracted from the endowments.

Finally, though we must have $b > 0$ as no KKO equilibrium exists, the limiting equilibrium as $b \rightarrow 0$ is well defined and has agent 2 consuming *all* of the good it has utility for. This demonstrates that while imposing a small minimum on endowments is not an economically viable substitute for resale, the limiting resale equilibrium is a well motivated alternative to imposing a small minimum on endowments.