Decentralized Convergence to Equilibrium Prices in Trading Networks

Edwin Lock[1], Benjamin Patrick Evans[2], Eleonora Kreacic[2], Sujay Bhatt[3], Alec Koppel[3], Sumitra Ganesh[3], Paul W. Goldberg[1]

[1]Department of Computer Science
University of Oxford
Parks Road
Oxford, OX1 3QD, UK
[2]JPMorgan AI Research
London, UK
[3]JPMorgan AI Research
New York, USA
{edwin.lock, paul.goldberg}@cs.ox.ac.uk
{benjamin.x.evans, eleonora.kreacic, sumitra.ganesh}@jpmorgan.com,
{sujay.bhatt, alec.koppel}@jpmchase.com

Abstract

We propose a decentralized market model in which agents can negotiate bilateral contracts. This builds on a similar, but centralized, model of trading networks introduced by Hatfield et al. in 2013. Prior work has established that fully-substitutable preferences guarantee the existence of competitive equilibria which can be centrally computed. Our motivation comes from the fact that prices in markets such as over-the-counter markets and used car markets arise from *decentralized* negotiation among agents, which has left open an important question as to whether equilibrium prices can emerge from agent-to-agent bilateral negotiations. We design a best response dynamic intended to capture such negotiations between market participants. We assume fully substitutable preferences for market participants. In this setting, we provide proofs of convergence for sparse markets (covering many real world markets of interest), and experimental results for more general cases, demonstrating that prices indeed reach equilibrium, quickly, via bilateral negotiations. Our best response dynamic, and its convergence behavior, forms an important first step in understanding how decentralized markets reach, and retain, equilibrium.

Extended version — https://arxiv.org/abs/2412.13972

1 Introduction

Price discovery is a central feature of modern markets. In many real-world markets, prices arise from a decentralized process governed by negotiations between market participants. By contrast, market outcomes are classically studied from the perspective of 'static' optimality criteria such as competitive equilibrium or stability, or viewed through the lens of centralized market procedures (such as auctions and clearing houses). In particular, this does not address the question of when, and how, desirable market outcomes can arise from *decentralized* market processes.

This paper introduces a 'best response' market dynamic that captures uncoordinated negotiations between neighbors in a network of market participants. We work in a general model of trading introduced by Hatfield et al. (2013), in which a network of heterogeneous agents can engage in bilateral trades for indivisible goods or services. Agents can act as buyer, sellers, or as traders involved in some trades as a buyer and in other trades as a seller. Our model allows for cycles of trades, and for 'horizontal' as well as 'vertical' (Ostrovsky 2008) network structures. The model also embeds well-known settings such as labor markets (Kelso and Crawford 1982), and general exchange economies with heterogeneous indivisible goods (Gul and Stacchetti 1999; Yang 2000), capturing a broad range of market configurations.

Like Hatfield et al., we focus primarily on agents with fully-substitutable (FS) preferences. Our central point of departure from Hatfield et al. (2013) is that each agent maintains an offer for each of their trades, so each trade has a buying *and* selling offer. In our market dynamic, agents take turns to update their offers in a manner that maximizes their utility. The dynamic ends in equilibrium if there is no agent who wishes to modify her offers any further (because her current offers already maximize her utility); at this point the buyer and seller of each trade agree to execute this trade iff their two offers coincide.

Trading networks have received increased attention in recent years (Morstyn, Teytelboym, and McCulloch 2018; Hatfield et al. 2019; Osogami, Wasserkrug, and Shamash 2023). Prior work has focused on establishing the existence of competitive equilibrium prices and stable outcomes in markets of increasing generality under the condition that agents have substitutes preferences (see, e.g., Kelso and Crawford 1982; Gul and Stacchetti 1999; Jagadeesan and Teytelboym 2021; Baldwin et al. 2023b). In the model of Hatfield et al. (2013), the authors show that competitive equilibra are guaranteed to exist if agents have fully-substitutable preferences. Moreover, an equilibrium is efficiently computable via an ascending-price algorithm. This algorithm can be run, e.g., by a central exchange to centrally compute prices if it has access to all agents' preferences.

But in many real world markets, such as over-the-counter financial markets and peer-to-peer markets, prices are not computed centrally, and preferences are not public. Instead, prices arise from uncoordinated bilateral negotiations of agents with their neighbors in a trading network. Such decentralized processes have been studied in the context of specialized two-sided markets; e.g., the labor markets (Chen, Fujishige, and Yang 2016a,b) and two-sided markets of buyers and sellers with unit demand for identical items (Assadi et al. 2017). While the decentralized market dynamics proposed in these works converge to competitive equilibrium, they are not immediately generalizable to markets on more general networks. Moreover, to the best of our knowledge, no similar results are known for market topologies beyond two-sided markets.

This leaves open the important research questions of formulating a decentralized market dynamic that captures the notion of negotiations between neighboring agents in general trading networks, and understanding the dynamic's (non-)convergence to an equilibrium.

Main Contributions. Extensive experiments suggest that our proposed market dynamic converges to equilibrium for any *market with fully substitutable agents (FSM)*. Section 3 develops theoretical results to support this observation. Our main technical contribution is a reduction from markets with arbitrarily many agents to 2-agent markets (Proposition 9). This reduction guarantees that FSMs with m-sparse networks converge (almost surely) to equilibrium iff FSMs with two agents and m trades converge (almost surely).[1]

As we show that 2-agent FSMs with one or two trades converge (Proposition 5), our reduction thus implies that markets with tree topologies and 2-sparse FSMs reach equilibrium under our market dynamic. We conjecture, based on our experimental evidence, that m-sparse FSMs converge for any $m > 2$. Surprisingly, we see that full substitutability is not required to guarantee convergence for markets with tree topologies; agents can have arbitrary quasilinear preferences (Theorem 10). By contrast, we give a counterexample of a two-agent market with a non-fully-substitutable agent for which the dynamic does not terminate.

Section 4 complements our theoretical results with multi-agent simulations to develop a better quantitative understanding of convergence in our market dynamic.[2] Specifically, we demonstrate the convergence, and speed of convergence, in markets with tree topologies, and move beyond the theoretical results to also demonstrate convergence in general (graph-based) markets. We explore numerous market configurations, including various market sizes and agent compositions, to analyze convergence and equilibrium properties. Our experiments demonstrate that equilibrium is reached much faster than the identified bounds. The experiments also highlight natural and desirable properties of our market dynamic, such as a general increase in welfare and satisfaction as we reach equilibrium, and the increase in utility as competition reduces (and vice versa). We also show how shocks propagate through the market, with knock-on effects before converging to a new equilibria, with the convergence speed governed by the severity of the shock.

The remainder of the paper is presented as follows. Section 2 proposes the model, Section 3 provides theoretical proofs, and Section 4 analyzes the simulation results. Conclusions and next steps are presented in Section 5.

[1] A network is *m-sparse* if every induced subgraph of the network can be divided into two or more disjoint components by removing at most m edges. See Definition 8 for details.

[2] For the code and implementation details, we refer to the extended version of this paper, available at https://arxiv.org/abs/2412.13972.

2 Model

A market $\mathcal{M} = (I, \Omega, v)$ consists of a finite set I of agents, a set Ω of possible bilateral trades and agent valuations $v = (v^i)_{i \in I}$. The trades Ω can represent any goods or services (such as the sale of a batch of coffee beans, an insurance contract, or a spectrum license), and are typically heterogeneous. Each trade ω has a single buyer $b(\omega) \in I$ and a single, distinct, seller $s(\omega) \in I$. Subsets of Ω are *bundles*. Ω can contain multiple trades with the same buyer and seller, and an agent can be a buyer for some trades and a seller for others. So Ω defines a directed multi-graph on vertices I in which every trade is an arc. Example 1 illustrates our model.

Example 1. The coffee supply chain consists of coffee farms, coffee roasters, coffee shops and supermarkets, forming the different types of agents. Roasters source raw beans from farms and process them to various varieties of blends, whilst coffee shops and supermarkets can source beans both from the roasters (processed blends) or farms (raw), representing the goods. A coffee shop/supermarket reaches out to farms and roasters in its network in order to obtain their asking price, and subsequently provides a set of buying offers (the trades). Figure 1 shows a concrete example. Notice the heterogeneous nature of goods the roaster trades: it acts as a buyer of raw product and a seller of various blends, as well as the heterogeneity of the supermarket and coffee shop which have different supplier networks and (private) preferences.

For any agent $i \in I$ and bundle $\Psi \subseteq \Omega$ of trades, Ψ_i denotes the trades in Ψ that involve agent i; and $\Psi_{i\to}$ and $\Psi_{i\leftarrow}$ respectively denote the agent's 'selling' and 'buying' trades in Ψ_i. For each agent i and trade ω, let $\chi^i_\omega = 1$ if i is the buyer of ω, $\chi^i_\omega = -1$ if i is the seller, and $\chi^i_\omega = 0$ otherwise.

Each agent i maintains an integer *offer* σ^i_ω for each of her trades $\omega \in \Omega_i$. The offer $\sigma^{b(\omega)}_\omega$ is an offer by the buyer $b(\omega)$ to *buy*, while $\sigma^{s(\omega)}_\omega$ is an offer by the seller $s(\omega)$ to *sell*. For any agent i and trade $\omega \in \Omega_i$, we write the offer of the other agent—its *counterpart*—as σ^{-i}_ω. Each agent interprets her counterparts' offers as *prices* $\boldsymbol{p} \in \mathbb{Z}^{\Omega_i}$ for trades Ω_i.

2.1 Preferences

As in Hatfield et al. (2013), agents have heterogeneous and private preferences. An agent's private *valuation function* v^i

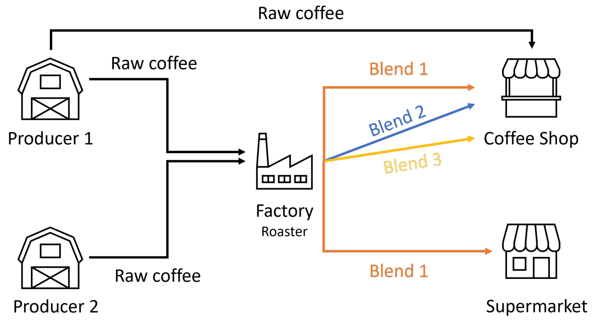

Figure 1: An illustration of a market with two producers (coffee farms), one intermediary (a roaster), and two buyers (coffee shop and supermarket).

maps every possible bundle $\Psi \subseteq \Omega_i$ of her trades to an integer value, and the empty bundle $\emptyset$ to 0.[3] The utility $u^i(\Psi, \boldsymbol{p})$ an agent has for bundle Ψ at prices $\boldsymbol{p} \in \mathbb{Z}^{\Omega_i}$ is the agent's value for Ψ plus the income from selling trades $\Psi_{i\rightarrow}$ minus the spending on buying trades $\Psi_{i\leftarrow}$:[4]

$$u^i(\Psi, \boldsymbol{p}) := v^i(\Psi) - \sum_{\omega \in \Psi} \chi^i_\omega p_\omega. \tag{1}$$

This gives rise to an agent's demand D^i consisting of the bundle $\Psi \subseteq \Omega_i$ that maximizes her utility at prices $\boldsymbol{p}$:[5]

$$D^i(\boldsymbol{p}) := \operatorname{argmax}_{\Psi \subseteq \Omega_i} u^i(\Psi, \boldsymbol{p}). \tag{2}$$

If an agent i is indifferent between multiple bundles at $\boldsymbol{p}$, we assume a consistent tie-breaking rule so that her demanded bundle is unique. In our experiments, we do this by perturbing v^i to achieve unique demand at any integer prices.

We focus on agents with fully-substitutable demands. An agent's demand is *fully-substitutable* if reducing the prices of some buying trades (weakly) decreases her demand for the remaining buying trades and (weakly) increases her demand of selling trades; and, conversely, raising the prices of some selling trades (weakly) decreases her demand for the remaining selling trades and (weakly) increases her demand of buying trades. For agents with only buying trades, or only selling trades, this definition coincides with the standard definition of substitutes introduced by Kelso and Crawford (1982) in the context of labor markets, also prevalent in the auction literature (Ausubel 2006; Baldwin et al. 2023a).

Definition 2. The demand of agent i is *fully-substitutable* if the uniquely demanded bundles Ψ, Ψ' at any prices $\boldsymbol{p}, \boldsymbol{p}' \in \mathbb{Z}^{\Omega_i}$ satisfy:

(i) $\Psi_{i\rightarrow} \subseteq \Psi'_{i\rightarrow}$ and $\{\omega \in \Psi'_{i\leftarrow} \mid p_\omega = p'_\omega\} \subseteq \Psi_{i\leftarrow}$ when $p_\omega = p'_\omega$ for $\omega \in \Omega_{i\rightarrow}$ and $p_\omega \geq p'_\omega$ for $\omega \in \Omega_{i\leftarrow}$;

(ii) $\Psi_{i\leftarrow} \subseteq \Psi'_{i\leftarrow}$ and $\{\omega \in \Psi'_{i\rightarrow} \mid p_\omega = p'_\omega\} \subseteq \Psi_{i\rightarrow}$ when $p_\omega = p'_\omega$ for $\omega \in \Omega_{i\leftarrow}$ and $p_\omega \leq p'_\omega$ for $\omega \in \Omega_{i\rightarrow}$.

We give examples of fully-substitutable buyers, sellers, and traders in the context of Example 1. Hatfield et al. (2015, 2019) provide further equivalent definitions of FS. We call a market with only fully-substitutable agents an *FSM*.

Example 3. The roaster from Fig. 1 interested in sourcing raw beans from any one of the two producers, but not both, can express her valuation v as $v(\emptyset) = 0$, $v(\{\omega_1\}) = v(\{\omega_2\}) > 0$ and $v(\{\omega_1, \omega_2\}) = -M$, where ω_1 and ω_2 denote the buying trades of the roaster and M is a sufficiently large negative number. Moreover, the roaster may have the technological constraint that she can only participate in a selling trade (e.g., ground coffee) if she also participates in the buying trade (e.g., coffee beans). She can implement this by assigning value $-M$ to any technologically infeasible bundle of trades, e.g., $v(\{\rho\}) = -M$, where ρ denotes any trade where the roaster acts as a seller.

From the perspective of the coffee shop, a reduction in the price of any blend would reduce its demand for the remaining blends (cf. Definition 2).

2.2 Market Dynamic

In our market dynamic, agents negotiate with their neighbors by modifying their own offers. Initially, the offers for trades are arbitrary, and all agents are marked as *unsatisfied*. The dynamic then selects one agent $i \in I$ at a time, uniformly at random, to update her offers by a discrete step size, $\varepsilon \in \mathbb{N}$.[6]

The agent i first observes her counterparts' offers for its trades Ω_i, which she interprets as the current prices $\boldsymbol{p} \in \mathbb{Z}^{\Omega_i}$ of these trades. She then determines the bundle $\Psi \in D^i(\boldsymbol{p})$ she demands at these prices, so that Ψ maximizes her utility $u^i(\cdot, \boldsymbol{p})$. She then *best responds* by updating her offers as follows. As the agent agrees to the prices of all trades Ψ, she sets her offers σ^i_ω for $\omega \in \Psi$ to match her counterparts' offers σ^{-i}_ω. For all trades $\omega \in \Omega_i \setminus \Psi$, which she does not demand, the agent sets her offer ω^i_ω to be ε lower than σ^{-i}_ω if ω is a buying trade, and ε higher than σ^{-i}_ω if ω is a selling trade. Setting her offers close to the trading partners' offers leaves room for negotiations while indicating her demand. Once an agent has best responded, she is marked as *satisfied* and all agents faced with a modified offer from agent i are marked as *unsatisfied*.

Definition 4. Suppose an agent i demands $\Phi \in D^i(\boldsymbol{p})$ at the prices $p_\omega = \sigma^{-i}_\omega$ corresponding to her counterpart's offer for each $\omega \in \Omega_i$. Her *best response* (BR) sets her new offers to

$$\sigma^i_\omega := \begin{cases} \sigma^{-i}_\omega & \text{if } \omega \in \Phi, \\ \sigma^{-i}_\omega - \chi^i_\omega \varepsilon & \text{else.} \end{cases} \tag{3}$$

The *state* $(U, \boldsymbol{\sigma})$ of the market dynamic after each BR is given by the set U of unsatisfied agents and the current of-

[3] Values for certain bundles can be negative, e.g. when they represent production costs for a seller or intermediary agent, as illustrated in Example 1.

[4] Heterogeneous preferences and quasilinear utilities are standard in the literature on substitutes markets and auctions (Hatfield et al. 2013; Kelso and Crawford 1982; Ausubel 2006).

[5] If an agent i is indifferent between multiple bundles at $\boldsymbol{p}$, we assume a consistent tie-breaking rule so that her demanded bundle is unique.

[6] Without loss of generality, we will assume $\varepsilon = 1$ throughout.

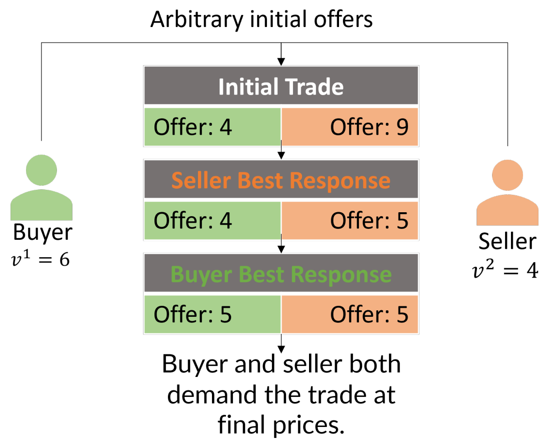

Figure 2: An example trade between two parties. Rather than trades having a single price, each trade has two offers, the buyer's offer and the seller's offer. Offers are initially arbitrary, and iteratively updated through best responses.

Algorithm 1 Market dynamic.

Initialize buying and selling offers $\boldsymbol{\sigma}$ for all trades.
Let $U = I$ be the set of unsatisfied agents.
while $U \neq \emptyset$ **do**
  Sample an agent $i \sim U$ uniformly at random.
  Let $\boldsymbol{p} \in \mathbb{Z}^{\Omega_i}$ with $p_\omega = \sigma_\omega^{-i}$ for all $\omega \in \Omega_i$ be the offers facing agent i.
  Determine the unique bundle $\Psi \in D^i(\boldsymbol{p})$ demanded by agent i at prices $\boldsymbol{p}$ [cf. (2)].
  Update agent i's offers to $\sigma_\omega^i = p_\omega$ for $\omega \in \Psi$ and $\sigma_\omega^i = p_\omega - \chi_\omega^i \varepsilon$ for $\omega \in \Omega_i \setminus \Psi$.
  Remove agent i from U.
  Add all agents $j \neq i$ facing modified offers to U.
return final offers $\boldsymbol{\sigma}$.

fers $\boldsymbol{\sigma}$. The dynamic terminates once all agents are satisfied, in which case we say that the market has reached *equilibrium*. Algorithm 1 describes the market dynamic in detail. At equilibrium, both agents i, j of each trade ω demand ω if $\sigma_\omega^i = \sigma_\omega^j$, and both don't demand ω if $\sigma_\omega^i \neq \sigma_\omega^j$.

As soon as every agent in the market has been selected at least once, it is straightforward that the two offers of any trade differ by 0 or ε, and in particular satisfy $\sigma_\omega^{b(\omega)} \leq \sigma_\omega^{s(\omega)} \leq \sigma_\omega^{b(\omega)} + \varepsilon$. We restrict our attention to this "main" phase of the market dynamic in our theoretical analysis.

We note that fully-substitutable utility functions $u^i(\cdot, \boldsymbol{p})$ are $M^\natural$-concave for any fixed prices $\boldsymbol{p}$ (cf. Hatfield et al. (2015, Section 4.5)). So agents can compute their demand at any prices $\boldsymbol{p}$ in polynomial time (Murota 2003, Chapter 10) to determine their best response.

3 The Emergence of Equilibria

In extensive numerical experiments, we observe that FSMs always converge to equilibrium. We now present theoretical results towards a proof that FSMs converge almost surely.[7] This qualitative result is complemented by our experiments in Section 4, which focus primarily on showing the convergence speed and additional welfare analysis. All formal statements and proofs can be found in the extended version. Let V be an upper bound on the absolute value of the agents' valuation functions and initial offers.

Any best response sequence for a market (I, Ω, v) in a given state is fully determined by the sequence $i_1, i_2, \ldots$ of agents $i_k \in I$ who are best-responding. In markets with two agents, the best response sequence for any given market and state is unique up to the starting agent, as the agents best respond in alternation. If the market consists of a single trade, then we see that it converges after at most $O(V)$ best responses. For two-agent markets with two trades, we show that full substitutability for both agents guarantees convergence to equilibrium after at most $O(V^2)$ best responses. We note that the definition of FS (Definition 2) only applies in the presence of two or more trades.

Proposition 5. *Markets consisting of a single trade converge to equilibrium after $O(V)$ BRs. Two-agent FSMs with two trades converge to equilibrium after $O(V^2)$ BRs.*

We prove the single-trade case by contradiction: suppose that the dynamic does not terminate. Then we observe that one agent always accepts her counterpart's offer for the trade ω, while the other always rejects. So the two agent's offers both increase, or both decrease, by ϵ after each round of best responses from both agents. But the buyer (seller) demands $\emptyset$ instead of the trade if the price is high (low) enough, a contradiction. In the two-trade case, we generalize this approach and use full substitutability to reach the required contradiction. We refer to the extended version for the full proof of Proposition 5.

By contrast, Example 6 presents a market with two trades between a substitutes buyer and a complements seller for which the market dynamic fails to terminate. This demonstrates that FS is a necessary condition in the maximal domain sense, as the market dynamic can cycle if one or more agents are not fully-substitutable. We conjecture that two-agent FSMs converge for any number m of trades after at most $f(m, V) = O(V^m)$ best responses, but that this will require new proof techniques.

Example 6. Consider a two-agent market with a buyer and a seller and valuations given by:

	$\emptyset$	$\{\omega\}$	$\{\varphi\}$	$\{\omega, \varphi\}$
buyer	0	8	9	$-\infty$
seller	0	-6	-7	-9

The buyer has substitutes preferences, and the seller has complementary preferences that violate Definition 2. If the buyer makes initial offers $\sigma_\omega^b = 4$ and $\sigma_\omega^b = 5$ for the two trades, and the market dynamic starts with a best response

[7] Recall that the randomness in the market arises from the choice of unsatisfied agent in each round of the market dynamic.

from the seller, then the dynamic cycles after two best responses from each agent. Assuming agents apply the lexicographic tie-breaking rule $\emptyset \prec \{\omega\} \prec \{\varphi\} \prec \{\omega, \varphi\}$ to decide their demanded bundle between two or more utility-maximizing bundles, the offers made by the alternating buyer and seller are:

	$\boldsymbol{\sigma}^b$	$\boldsymbol{\sigma}^s$	$\boldsymbol{\sigma}^b$	$\boldsymbol{\sigma}^s$	$\boldsymbol{\sigma}^b$
ω	4	5	5	5	4
φ	5	6	5	5	5

Conjecture 7. Two-agent FSMs with any number of trades converge to equilibrium.

We next turn to markets with arbitrarily many agents. Our main theoretical result is to reduce the problem of convergence for markets with any number of agents to the case of two-agent markets. For this, we categorize markets by the 'sparsity' of their topologies.

Definition 8. A market is *m-sparse* if every subgraph of its underlying graph admits a cut of size at most m. (A *cut of size m* is a partition of a graph's vertices into two vertex sets with m crossing edges between the two sets.)

The topology of a 1-sparse market is a forest. Moreover, markets with agents who are each involved in at most m trades are trivially m-sparse.

In Proposition 9, we show that the convergence of two-agent markets with one trade implies that any 1-sparse market admits a best response sequence after which the market terminates. Surprisingly, this result holds even if the agents are not fully-substitutable. For m-sparse markets with $m \geq 2$, Example 6 motivates us to focus on the setting with full substitutability. We establish that every FSM admits a best response sequence after which the dynamic terminates if and only if two-agent FSMs converge.

Proposition 9.

(i) Every market with a forest topology (even when agents are not FS), in any state, admits a terminating best response sequence iff two-agent markets with one trade converge.

(ii) Every m-sparse FSM, in any state, admits a terminating best response sequence iff two-agent FSMs with m trades converge.

We prove Proposition 9 in two steps. Fix a market $\mathcal{M} = (I, \Omega, v)$ in state $(U^0, \boldsymbol{\sigma}^0)$. We partition the set of agents I of the market into two non-empty subsets, J^1 and J^2, and let $\Omega^{1,2}$ be the trades between these agent sets. This is illustrated in Fig. 3, in which $\Omega^{1,2} = \{\omega, \varphi, \psi\}$. For notational convenience, define $J^k := J^{k \bmod 2}$ for $k \geq 1$.

In the first step, we describe a procedure to construct a best response sequence $S = S^1 S^2 \ldots$ consisting of subsequences S^k. Each S^k contains only agents from J^k, and *all* agents in J^k are satisfied after applying sequence $S^1 \ldots S^k$. Our second step is then to prove that S is finite and the market terminates after applying S. We now sketch out each step in more detail; the full proof is given in the extended version.

The subsequences S^k are constructed recursively. Let $(U^{k-1}, \boldsymbol{\sigma}^{k-1})$ be the state of the market after applying sequence $S^1 \ldots S^{k-1}$. If all agents in the market are satisfied after this sequence, we are done and define $S = S^1 \ldots, S^{k-1}$. Otherwise, to find S^k, we restrict the market $\mathcal{M}$ in state $(U^{k-1}, \boldsymbol{\sigma}^{k-1})$ to agents J^k and modify the valuations of the agents in J^k who share a trade with agents outside J^k. The modified valuations endow agents with the opportunity to execute their trades in $\Omega^{1,2}$ at prices set to their counterparts' offers, even though these are no longer included in the restricted market. This valuation transformation was first described in Hatfield et al. (2015), and preserves full substitutability. Intuitively, this means that the agents J^k best respond identically in the original and restricted markets. By induction, the restricted market admits a terminating best response sequence S^k. As all agents in the restricted market are satisfied after S^k, all agents in J^k are also satisfied after applying S^k to the original market in state $(U^{k-1}, \boldsymbol{\sigma}^{k-1})$. Figure 3 (b) shows the restricted market in which the three agents of J^k involved in the grayed-out trades $\Omega^{1,2}$ have valuations that endogenise these trades.

In order to show that S is finite, we consider the market $\widetilde{\mathcal{M}}$ obtained by merging agent sets J^1 and J^2 into two agents, 1 and 2. The trades of $\widetilde{\mathcal{M}}$ consist of $\Omega^{1,2}$ with endpoints changed in the natural way to agents 1 and 2, and the initial market state is modified similarly. Moreover, the valuation of merged agent k is defined such that she demands the same trades $\Omega^{1,2}$ as the agents J^k in the original market do, whenever the counterparts' offers for these trades are the same in both markets, preserving FS (Hatfield et al. 2015). As $\widetilde{\mathcal{M}}$ has two agents, it admits a unique best response sequence starting with agent 1. We show that because all agents J^k in the original market are satisfied after sequence $S^1 \ldots, S^k$, the offers $\boldsymbol{\sigma}^k$ for trades $\Omega^{1,2}$ after this sequence are the same as the offers after the kth best response in the merged market. This implies that the length of the best response sequence for $\widetilde{\mathcal{M}}$ is the number of subsequences of S. As two-agent markets converge by assumption, S must thus converge after finitely many best responses, concluding the proof sketch.

Our market dynamic selects unsatisfied agents uniformly at random, so it will eventually select a terminating best response sequence if one exists. So Propositions 5 and 9 imply that 1-sparse markets (even with non-FS agents) and 2-sparse FSMs converge almost surely. Similarly, under the assumption that Conjecture 7 holds, every FSM converges almost surely.

Theorem 10.

(i) Every 1-sparse market converges almost surely to equilibrium, even when agents are not FS.

(ii) Every 2-sparse FSM converges almost surely to equilibrium.

(iii) Every FSM converges almost surely to equilibrium if Conjecture 7 holds.

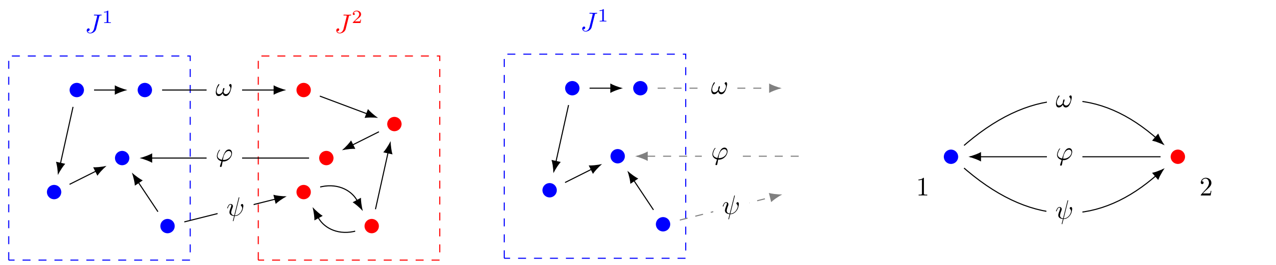

(a) A 3-sparse market (b) The market restricted to J^1 (c) The market after merging J^1 and J^2.

Figure 3: Illustrating the proof of Proposition 9.

4 Experiments

In the previous section, we proved theoretical convergence guarantees for 1-sparse and 2-sparse FSMs, and conjectured that all FSMs converge. We now explore the market dynamic computationally, analyzing paths and time to convergence under various market settings. Our primary objective for running experiments is to support the theoretical claims of the previous section and study the rate of convergence to equilibrium. Additionally, we analyze the resulting welfare based on agent composition, and the impact of exogenous shocks on the resulting equilibria and convergence.

Experiment Configuration. We assume one kind of good, and that sellers and buyers have unit supply/demand for this good, i.e., they want to buy/sell at most one item of this good[8]. Each buyer and seller's value $c_i \in C$ for an item of this good is drawn uniformly at random from a predetermined set of integers, $C = \{1, \ldots, 100\}$. For the seller, the value can be understood as the production cost, and for a buyer, their perceived fundamental value of the good. The valuation function of a buyer i is given by $v^i(\emptyset) = 0$, $v^i(\Psi) = c_i$ if $|\Psi \cap \Omega_i| = 1$ and $v^i(\Psi) = -\infty$ otherwise. For sellers, we have $v^i(\emptyset) = 0$, $v^i(\Psi) = -c_i$ if $|\Psi \cap \Omega_i| = 1$ and $v^i(\Psi) = -\infty$ otherwise. Intermediaries cannot sell more than they buy, and have no value for retaining items, so their valuation is $v^i(\Psi) = 0$ if Ψ_i contains the same number of buying and selling trades, and $v^i(\Psi) = -\infty$ otherwise. Environments are configured in Phantom (Ardon et al. 2023).

Market Topologies. We analyse several FSM topologies, ranging from bipartite networks of buyers and sellers to general networks with buyers, sellers, and intermediaries.

BS: Bipartite networks consisting of buyers and sellers. Buyers only trade with sellers and vice versa.

BIS: Networks with buyers, sellers, and intermediaries. Buyers and sellers only trade through intermediaries, and intermediaries do not trade with other intermediaries.

General: Networks with buyers, sellers, and intermediaries. Buyers and sellers transact through intermediaries or directly. Additionally, intermediaries may trade with other intermediaries, and cycles are permitted.

Examples of these topologies, and their construction processes, are described and visualized in the extended version. The BS topology models direct trade between buyers and sellers, e.g., direct-to-consumer or business-to-consumer markets. The other topologies, which include intermediaries, can model a number of more complex markets, such as used car markets with buyers and sellers interacting through dealers (Hatfield et al. 2013), energy markets with consumers and generators mediated through suppliers (Morstyn, Teytelboym, and McCulloch 2018), or even general supply chains (Ostrovsky 2008).

4.1 Experimental Results

The convergence rates across the network structures with different market *sizes* (number of market participants N), and *compositions* (the proportion of buyers, sellers, and intermediaries) are analyzed. Each configuration is run 100 times on a standard personal computer[9], with the mean $\pm$ standard deviation presented.

BS Networks. For each buyer/seller pair i, j in a BS network, a potential trade between i and j is added with some probability r (here, $r = 0.1$).

Convergence paths are visualized in Fig. 4a. Increasing the number of agents increases the time to convergence, as there is more room for bargaining/trade among the agents (e.g., a buyer has access to more sellers, so has more negotiation opportunity). Importantly, in each case, we see convergence (the satisfied rate converges to 1). Convergence rates across N are visualized in the extended version, Fig. 9.

Varying the proportion of buyers in the market (Fig. 5), we recover the expected economic welfare behavior that as the proportion of buyers increases (decreases), their welfare, measured by the average utility, decreases (increases), due to market competition, resulting in reduced consumer surplus (increased producer surplus).

BIS Networks. Introducing intermediaries into the network alters the convergence rates as transactions must happen through an intermediary. We explore different numbers of intermediaries, again, with a fixed connectivity rate of $r = 0.1$, where for the possible pairs i, j at least one of i

[8] Note that neither of these are requirements of the market setting, but allow for a clear evaluation.

[9] Run with Python 3 on a 2022 Macbook Air M2 with 8GB RAM.

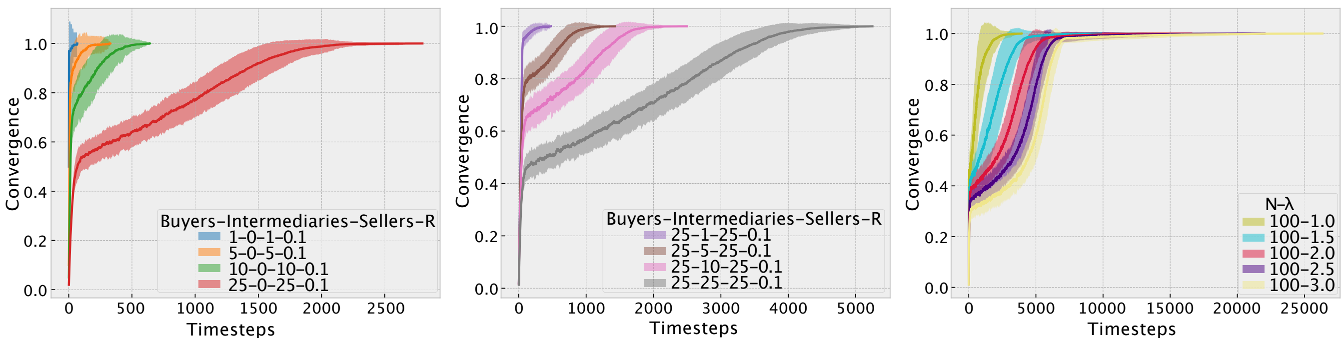

(a) BS (b) BIS (c) General

Figure 4: Convergence paths by market topology. The x-axis shows the number of iterations (e.g., number of best response updates), and the y-axis shows the proportion of *satisfied* agents.

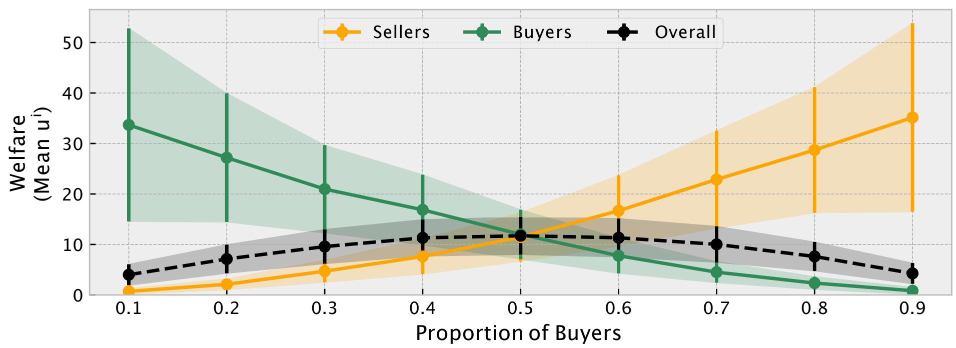

Figure 5: Social welfare (mean u_i) by agent composition in the BS network. The x-axis varies the proportion of buyers, and the y-axis displays the resulting welfare.

or j must be an intermediary (e.g., there is no direct connection between buyers and sellers).

Convergence rates across an increasing number of intermediaries are visualized in Fig. 4b. In each case, the market converges to an equilibrium, as predicted by the theory. Increasing the number of intermediaries increases the time to convergence, as there are more opportunities for trade. When the number of intermediaries is low, buyers/sellers have little room for negotiation, so convergence is more rapid. The more intermediaries in the market, the lower their utilities are due to the competition, which results in higher utilities for the buyers and sellers.

General Networks. We now consider general trading networks. For these experiments, we generate Erdős-Rényi networks with connectivity rates $r = \frac{\lambda}{N}$, where $N = 100$ and $\lambda = \{1, 1.5, 2, 2.5, 3\}$.

Convergence rates across λ are visualized in Fig. 4c. Despite featuring cycles, these trading networks still converge, at a rate approximately linear with λ (see Fig. 9 in extended version). While we do not yet have theoretical guarantees of convergence under this setting, these results help to show that experimentally these general topologies also converge, opening a promising line of future research building on Theorem 10. The welfare of the agents does not change drastically across λ, but generally, as λ increases the proportion of intermediaries in the network increases, so the welfare for buyers/sellers increases, and the welfare for intermediaries decreases (see extended version).

4.2 Shocks

To understand how exogenous shocks impact the market, shocks are applied to an agent's value c_i, resampling to a new value in C. Shocks have a specific size s, which controls the variation from the original valuation, i.e., a shock size of 0.1 means the updated $c_{i,t}$ will be in the range $[c_{i,t-1} * 0.9, c_{i,t-1} * 1.1]$ (clipped to the values in C). Shocks are applied to different proportions of agents in the network to analyze their propagation. Shocks only occur after the dynamics have converged to understand how these perturbations affect an otherwise converged system.

The results for the BS network are visualized in Fig. 6. We look at two key quantities: the proportion of additional (non-shocked) agents affected (Fig. 6a), e.g., due to knock-on effects of one borrower updating their valuation, and the normalized time taken to reconverge (Fig. 6b).

Propagation. When considering the propagation of shocks (Fig. 6a), a knock-on effect is observed in all cases, with

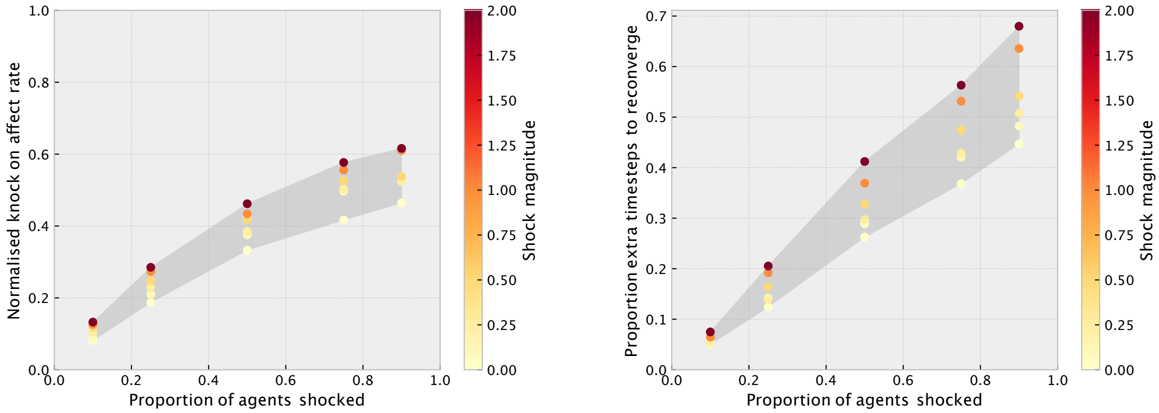

(a) Propagation

(b) Reconvergence speed

Figure 6: The impact of exogenous shocks in a BS network. Fig. 6a shows the propagation of shocks measured as the additional proportion of unsatisfied agents. Fig. 6b shows the normalized reconvergence speed (as a proportion of the original convergence speed). Each ∘ represents the mean impact for a shock size (darker color = larger magnitude). The filled region shows the effect range across these shock sizes.

additional non-shocked agents becoming affected (i.e., becoming unsatisfied and needing to update their prices). For example, when applying a shock to 25% of the agents, approximately $50\%+$ of the agents become impacted and must update their responses. Generally speaking, the larger the shock magnitude, the more propagation throughout the network is observed, agreeing with real world insights on shock propagation's, for example in production networks (Huneeus 2018) and global trade-investments (Starnini, Boguñá, and Serrano 2019), where the magnitude of the shock determines its impact on the overall system.

Reconvergence. The additional time to converge (Fig. 6b) increases approximately linearly with the proportion of shocked agents and the shock magnitude, showing the disruption effects. As expected by the theory, irrespective of the shock size, the dynamics eventually settle down and reconverge to a new equilibrium. This reconvergence is significantly faster than the original convergence speed (Fig. 6b), essentially serving as a warm start to the convergence process, which is particularly true for small shock magnitudes.

4.3 Results Summary

These experimental results support and extend the theoretical guarantees in Section 3, by simulating the market dynamic for various trading networks and establishing consistent convergence across market topologies. The best response sequences constructed in the proof of Proposition 9 can be exponentially long. By contrast, we show experimentally that these markets converge rapidly. Additionally, we demonstrate that more complex network structures also still converge under the proposed market dynamic, opening a promising line of future analysis extending Proposition 9. The experimental analysis provided insights into the resulting welfare and shock impacts across various market topologies, expanding the understanding of such market processes.

5 Discussion and Conclusion

We formulate a decentralized market dynamic in a heterogeneous network of trading agents. Our dynamic captures iterative negotiation among the agents, by allowing buyers (sellers) to refuse trades via making counteroffers that are slightly lower (higher) than the price being offered. We demonstrate theoretically and experimentally how prices converge to an equilibrium via an uncoordinated negotiation (from arbitrary initial prices), providing the first such analysis in general market settings.

Theoretically, we develop a reduction that ensures convergence of many-agent markets iff two-agent markets converge. We apply this reduction to prove that our market dynamic reaches equilibrium for markets with tree topologies, as well as 2-sparse FSMs. Experimentally, we provide empirical evidence and extensions to the convergence guarantees and convergence process, demonstrating that fully-substitutable markets converge to an equilibrium and do so in a manner significantly faster than suggested by the proofs in our theoretical work. Additionally, we highlight several natural and desirable features of the market, e.g., market reactions to exogenous shocks and the welfare impact on agents based on market compositions.

In addition to proving our Conjecture 7 about the convergence of two-agent markets with many trades, future work could address more sophisticated choices of counteroffers that may lead to faster convergence to equilibrium, e.g., modifying the step size ε in the spirit of adaptive stochastic gradient based optimization algorithms (Kingma and Ba 2015; Duchi, Hazan, and Singer 2011). Additionally, the agents in our market dynamic optimize their immediate utility, but future iterations could also model strategic traders (Vadori et al. 2024) who consider the long-term impact of their offers.

Acknowledgments

Goldberg and Lock were supported by a JP Morgan faculty fellowship and EPSRC grant EP/X040461/1.

This paper was prepared for informational purposes in part by the Artificial Intelligence Research group of JPMorgan Chase & Co. and its affiliates ("JP Morgan") and is not a product of the Research Department of JP Morgan. JP Morgan makes no representation and warranty whatsoever and disclaims all liability, for the completeness, accuracy or reliability of the information contained herein. This document is not intended as investment research or investment advice, or a recommendation, offer or solicitation for the purchase or sale of any security, financial instrument, financial product or service, or to be used in any way for evaluating the merits of participating in any transaction, and shall not constitute a solicitation under any jurisdiction or to any person, if such solicitation under such jurisdiction or to such person would be unlawful.

References

Ardon, L.; Vann, J.; Garg, D.; Spooner, T.; and Ganesh, S. 2023. Phantom-A RL-driven Multi-Agent Framework to Model Complex Systems. In *Proceedings of the 2023 International Conference on Autonomous Agents and Multiagent Systems*, 2742–2744.

Assadi, S.; Khanna, S.; Li, Y.; and Vohra, R. 2017. Fast convergence in the double oral auction. *ACM Transactions on Economics and Computation (TEAC)*, 5(4): 1–18.

Ausubel, L. M. 2006. An Efficient Dynamic Auction for Heterogeneous Commodities. *Am. Econ. Rev.*, 96(3): 602–629.

Baldwin, E.; Goldberg, P.; Klemperer, P.; and Lock, E. 2023a. Solving Strong-Substitutes Product-Mix Auctions. *Mathematics of Operations Research*.

Baldwin, E.; Jagadeesan, R.; Klemperer, P.; and Teytelboym, A. 2023b. The Equilibrium Existence Duality. *Journal of Political Economy*, 131(6): 1440–1476.

Chen, B.; Fujishige, S.; and Yang, Z. 2016a. Random decentralized market processes for stable job matchings with competitive salaries. *Journal of Economic Theory*, 165: 25–36.

Chen, B.; Fujishige, S.; and Yang, Z. 2016b. Random decentralized market processes for stable job matchings with competitive salaries. *Journal of Economic Theory*, 165: 25–36.

Duchi, J.; Hazan, E.; and Singer, Y. 2011. Adaptive subgradient methods for online learning and stochastic optimization. *Journal of Machine Learning Research*, 12(7).

Gul, F.; and Stacchetti, E. 1999. Walrasian equilibrium with gross substitutes. *Journal of Economic theory*, 87(1): 95–124.

Hatfield, J. W.; Kominers, S. D.; Nichifor, A.; Ostrovsky, M.; and Westkamp, A. 2013. Stability and competitive equilibrium in trading networks. *Journal of Political Economy*, 121(5): 966–1005.

Hatfield, J. W.; Kominers, S. D.; Nichifor, A.; Ostrovsky, M.; and Westkamp, A. 2015. Full Substitutability in Trading Networks. In *Proceedings of the Sixteenth ACM Conference on Economics and Computation*, EC '15, 39–40. New York, NY, USA: Association for Computing Machinery. ISBN 9781450334105.

Hatfield, J. W.; Kominers, S. D.; Nichifor, A.; Ostrovsky, M.; and Westkamp, A. 2019. Full substitutability. *Theoretical Economics*, 14(4): 1535–1590.

Huneeus, F. 2018. Production network dynamics and the propagation of shocks. *Graduate thesis, Princeton University, Princeton, NJ*, 52.

Jagadeesan, R.; and Teytelboym, A. 2021. Matching and money. In *Proceedings of the 22nd ACM Conference on Economics and Computation*, 634–634.

Kelso, A. S.; and Crawford, V. P. 1982. Job matching, coalition formation, and gross substitutes. *Econometrica: Journal of the Econometric Society*, 1483–1504.

Kingma, D. P.; and Ba, J. 2015. Adam: A Method for Stochastic Optimization. In Bengio, Y.; and LeCun, Y., eds., *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*.

Morstyn, T.; Teytelboym, A.; and McCulloch, M. D. 2018. Bilateral contract networks for peer-to-peer energy trading. *IEEE Transactions on Smart Grid*, 10(2): 2026–2035.

Murota, K. 2003. *Discrete Convex Analysis: Monographs on Discrete Mathematics and Applications 10*. USA: Society for Industrial and Applied Mathematics. ISBN 0898715407.

Osogami, T.; Wasserkrug, S.; and Shamash, E. S. 2023. Learning Efficient Truthful Mechanisms for Trading Networks. In *Proceedings of the Thirty-Second International Joint Conference on Artificial Intelligence, IJCAI-23*, 2862–2869.

Ostrovsky, M. 2008. Stability in supply chain networks. *American Economic Review*, 98(3): 897–923.

Starnini, M.; Boguñá, M.; and Serrano, M. Á. 2019. The interconnected wealth of nations: Shock propagation on global trade-investment multiplex networks. *Scientific Reports*, 9(1): 13079.

Vadori, N.; Ardon, L.; Ganesh, S.; Spooner, T.; Amrouni, S.; Vann, J.; Xu, M.; Zheng, Z.; Balch, T.; and Veloso, M. 2024. Towards multi-agent reinforcement learning-driven over-the-counter market simulations. *Mathematical Finance*, 34(2): 262–347.

Yang, Z. 2000. Equilibrium in an exchange economy with multiple indivisible commodities and money. *Journal of Mathematical Economics*, 33(3): 353–365.