Liquidity Risk in Sequential Trading Networks*

Shachar Kariv† Maciej H. Kotowski‡ C. Matthew Leister§

October 2, 2016

Abstract

This paper develops a model of intermediated exchange with budget-constrained traders who are embedded in a trading network. An experimental investigation confirms the theory's baseline predictions. Traders adopt monotone strategies with higher-budget intermediaries offering to pay more for tradable assets. Traders closer to the final consumer in the network experience systematically greater payoffs due to lessened strategic uncertainty. While private budget constraints inject uncertainty into the trading environment, they also serve as a behavioral speed-bump, preventing traders from experiencing excessive losses due to overbidding.

Keywords: Experiment, Economic Networks, Intermediation, Liquidity, Auctions, Strategic Uncertainty
JEL: L14, C91, D85, D44, G10

*We thank Douglas Gale for detailed comments and suggestions. We are also grateful to Larry Blume, Gary Charness, Syngjoo Choi, Andrea Galeotti, Edoardo Gallo, Sanjeev Goyal, Tom Palfrey, Brian Rogers, and Bill Zame for helpful discussions and comments. The paper has also benefited from suggestions by the participants of seminars at several universities. This research was supported by the UC Berkeley Experimental Social Science Laboratory (Xlab), and the experiments were funded by the Multidisciplinary University Research Initiatives (MURI) Program—Multi-Layers and Multi-Resolution Networks of Interacting Agents in Adversarial Environments.

†Department of Economics, University of California, Berkeley, 530 Evans Hall #3880, Berkeley, CA 94720, United States. E-mail: kariv@berkeley.edu.

‡John F. Kennedy School of Government, Harvard University, 79 JFK Street, Cambridge MA 02138, United States. E-mail: maciej_kotowski@hks.harvard.edu.

§Department of Economics, Monash University, 900 Dandenong Road, Caulfield East VIC 3145, Australia. E-mail: matthew.leister@monash.edu.

1 Introduction

Liquidity lubricates trade. At each moment in time, a trader needs to have adequate financial resources—either in the form of cash on hand or credit—to pay for the assets and goods that he wishes to buy. When liquidity is ample, markets work effectively and exchange flows freely. When agents' financial capacity dwindles, risk is amplified, counter-parties may not get paid, and trade grinds to a halt, much like an engine in need of oil.

The risks posed by a particular trader's lack of financial capacity would be small if he was but one agent in a very large pool of anonymous market participants. A financially-sound competitor, the argument goes, could easily step-in and ensure that trade continues. Most economic transactions, however, depart significantly from this large-market ideal, common to economic analysis. Rather, decentralized markets often rely on a few established trading relationships that form a network (Jackson, 2008, Chapter 10). Traders serve as intermediaries or middlemen between a good's producer and its ultimate consumer (Rubinstein and Wolinsky, 1987). A producer sells his output to a few reliable wholesalers. Each wholesaler supplies a few retailers. Each retailer has a particular clientele. This network is often incomplete with intermediaries knowing or trusting only a few others in the market.

The interaction between traders' financial risk and market imperfections summarized by a network naturally raises several open questions:

1. How do intermediary traders account for others' (possibly) limited financial capacity?
2. What price dynamics might we observe as goods are bought, resold, and ultimately consumed?
3. What are the distributional consequences of the economy's network structure? Are some traders systematically advantaged because of their position in the market?

Despite the situation's ubiquity in both markets for financial assets and for physical goods, precise answers to these questions are far from obvious.

In this paper we are the first to investigate the interaction between an economy's extent of decentralization and the private financial capacity of its participants. A profound obstacle in answering the above questions is that many key variables—the agents' preferences, their private financial constraints, their strategies—are not readily observable. Thus, we pursue complementary tracks involving theoretical and experimental analysis. We develop a tractable and testable model that we take to the laboratory where we can directly observe and control the private variables guiding behavior.

Our economic model is designed to be simple, though rich enough to adequately probe the above questions. There is a producer of a good or asset (the "seller") and a final consumer (the "buyer"). We assume that the seller and the buyer cannot trade directly. Instead, the asset must pass through a sequence of intermediary traders en route from the seller to the buyer. A network defines the relationships among the traders and is common knowledge. If trader i is linked to trader j, then these agents can trade; otherwise, transactions between them cannot occur. Traders, whose motives are purely speculative, buy and resell the asset facilitating its eventual passage along links in the network from the seller to the buyer. The price in each transaction is determined by a first-price auction, a well-understood and familiar pricing mechanism.

The key feature of our model is that each trader is liquidity or budget constrained and has limited funds to finance his trading activity. As true in practice, these constraints are private information and inject considerable risk into the market. A trader may be forced to sell the asset at a loss if his neighbors' financial constraints are binding. A successful trader must anticipate the funds available to his immediate counter-parties and be mindful of similar constraints present elsewhere in the economy. This compounding of uncertainty through the network distinguishes our analysis from that of Gale and Kariv (2009), who assume a perfect information setting. In equilibrium, intermediary traders in our economy adopt subtle bidding strategies where bids are shaded on account of others' liquidity risks. Prices rise systematically as the asset nears the final buyer, and traders closer to the buyer benefit directly from the reduction of strategic uncertainty in order to earn higher trading profits.

Our model's simplicity ensures that it can be readily tested empirically in a laboratory setting. Our experiments confirm the key conclusions from our theoretical analysis. Traders with high endowments strategically reduce their bid relative to their budget. Across networks, intermediaries closer to the final buyer adopt uniformly more aggressive bidding strategies. Consequently, resale prices rise as the tradable asset approaches the buyer and intermediaries closer to the buyer tend to earn higher trading profits. Interestingly, our experiment also identifies a practical disciplinary role played by liquidity and budget constraints. We observe that traders who are flush with resources tend to experience a mild decline in profits relative to others and to the equilibrium benchmark. While a lax budget constraint allows a trader greater freedom to pursue his goals, it also allows him to err and overbid with greater frequency. Contrary to standard theory, removing this constraint can be a mixed blessing in practice as the latter effect may dominate.

This paper is organized as follows. In Section 2 we briefly review the related literature. We link our study to both theoretical and experimental studies of networked markets and auctions. In Section 3 we formally introduce our baseline model and use it to derive several implications concerning intermediary behavior, price dynamics, and welfare. In Section 4 we describe the experimental investigation. We carefully explain how the experimental market captures the key features of our theoretical model. In Section 6 we analyze our experiment's results. Section 6 concludes. Proofs are gathered in Section 7. The paper also includes several data and technical online appendices for the interested reader.

2 Related Literature

This paper straddles two distinct, though complementary literatures. From a theoretical point of view, our study is closely related to recent examinations of sequential trade within networked economies. This literature builds upon Kranton and Minehart (2001) by incorporating market intermediaries into the trading process. By assumption, buyers and sellers must interact through intermediaries, who facilitate trade while respecting their network of relationships. Recent contributions to this literature include Gale and Kariv (2007), Blume et al. (2009), Gale and Kariv (2009), Kotowski and Leister (2014), Choi et al. (2014), Manea (2014), and Condorelli et al. (2016). Condorelli and Galeotti (2014) provide a recent survey. Our study complements this growing literature by proposing a new model of intermediated exchange in the presence of private information about trading ability.

Since the pioneering studies by Smith (1962, 1965), a large literature in experimental economics has investigated the operation of markets under various conditions. Our analysis is closest to that of Gale and Kariv (2009). They study a similar class of trading economies, albeit in a complete-information setting. As expected in their complete-information setting, Gale and Kariv (2009) show that prices converge to the asset's value and all intermediation rents are quickly competed away. Going beyond Gale and Kariv (2009), we study an economy where intermediaries hold private information concerning their trading ability, which we model as an endowment of funds to bankroll transactions. We show that private information has considerable implications for traders' strategic calculus, the resulting price dynamics, and intermediaries' welfare. For example, our theory predicts increasing price sequences. Moreover, intermediaries' strategies allow them to capture positive intermediation rents and their welfare is systematically affected by their position in the trading network. Our theory and experiment let us describe and test both of these features.

Other network-based trading games are studied by Charness et al. (2007) and Judd and Kearns (2008). The former considers bargaining while the latter considers a trading process with continuously-updated limit-order books. Recently, Choi et al. (2014) have examined intermediated exchange theoretically and in the laboratory focusing on a posted-price mechanism. Their experiment emphasized the importance of coordination among market intermediaries and the role of so-called "critical traders." Critical traders enjoy monopoly-like positions in a network allowing them to extract considerable intermediation rents. Kosfeld (2004) provides a survey of other laboratory experiments involving networked economies.

Though our focus is on trading dynamics in a networked economy, our study also contributes to the experimental auction literature. Our experiments incorporate a direct test of Che and Gale's (1996) model of a common-value, first-price auction with private budget constraints.[1] To our knowledge, Che and Gale's (1996) model has not been subject to laboratory verification before. Other experimental investigations of auctions with private budget constraints include Kotowski (2010) and Ausubel et al. (2013), albeit they consider environments different from ours. Pitchik and Schotter (1988) examine sequential auctions with budget constraints theoretically and experimentally; however, they consider a setting where multiple goods are sold in sequence. In our case, the same good is sold and resold among a sequence of distinct agents. Comprehensive surveys of the experimental auction literature are provided by Kagel (1995) and Kagel and Levin (2014).

3 The Model

In this section we introduce a baseline model that will inform our experimental analysis. We discuss modeling choices and model interpretation in the following subsection. Thereafter, we characterize the economy's equilibrium to derive testable implications that we will take to the laboratory.

Consider an economy where trading possibilities are summarized by a directed graph. Agents are nodes and a directed link from agent i to agent j indicates that i can sell to j. The network structure is exogenous and common knowledge. Though our theory generalizes, we focus on a tractable class of networks originally analyzed by Gale and Kariv (2009). Choi et al. (2014) call such networks "multipartite networks" and we borrow their terminology

[1] Che and Gale (1998) and Kotowski (2016) examine first-price auctions in a private-value setting where the agents' valuations and budgets are private information. Equilibria in such environments can be very complicated due to the type-space's multi-dimensionality. Bobkova (2016) generalizes the model of Che and Gale (1996) to the case of asymmetric budget distributions.

when convenient. Agents are arranged in rows $0, 1, \ldots, R+1$. One agent is the buyer (B) and one agent is the seller (S). The buyer is located in row 0 and the seller is located in row $R+1$. The buyer is offering to pay $v > 0$ for an asset or good that is created by the seller at zero cost (a normalization).[2] There exist gains from transferring the asset from the seller to the buyer. However, these agents are not linked directly. Instead, there is a set of intermediary traders who may buy and sell the asset. In our model, they inhabit rows $1, \ldots, R$. Rows are numbered according to their network distance from the buyer.[3] Traders, who are risk neutral, do not value the asset per se. Rather they seek to earn profits by facilitating trades in the network. As usual, a trader earns profits if he can "buy low" and (re)sell "high." The network constrains the set of feasible trades as follows. An agent in row r can purchase the asset from an agent in row $r+1$ and can sell the asset to an agent in row $r-1$.[4]

Figure 1 illustrates a family of admissible trading networks. These particular trading networks feature in our experimental investigation and involve one, two, or three rows of intermediaries. Intuitively, as the distance between the buyer and the seller increases, the market's operation involves a more complex chain of intermediary transactions. Throughout our analysis, we assume that there is a common number $N \geq 2$ of intermediary traders in each row. Extending our results to allow for a different number of intermediaries in each row is straightforward and does not change our qualitative conclusions. In Figure 1 each row has three intermediaries.

All trade occurs via sequential auctions according to the following protocol. When an agent in row $r+1$ holds the asset, he organizes a first-price, sealed bid auction to sell it. Agents in row r participate in the auction by submitting bids. The highest bidder is awarded the asset and makes a payment equal to his bid. Ties among highest bidders are resolved with a uniform randomization. This process continues until the asset reaches the buyer, who pays v for the asset.

As described, the model above broadly parallels that of Gale and Kariv (2009).[5] Substan-

[2] Following Gale and Kariv (2009), we call the traded good an "asset." Depending on the application, it may represent a financial product or a physical item.

[3] The row numbering follows Kotowski and Leister (2014), which is opposite to the convention followed by Gale and Kariv (2009). Numbering rows from the "end" simplifies the equilibrium characterization.

[4] A richer model may additionally introduce links between traders *within* a row. It can be shown that these additional links will not impact our economy's equilibrium. For simplicity, we omit them from our theoretical and experimental analysis.

[5] Gale and Kariv (2009) adopt a different price setting mechanism involving reserve, or ask, prices. In their baseline treatment the transaction price was an average of the bid and the ask price. They report results from a control treatment where transaction prices equaled only the bid. These are similar to their

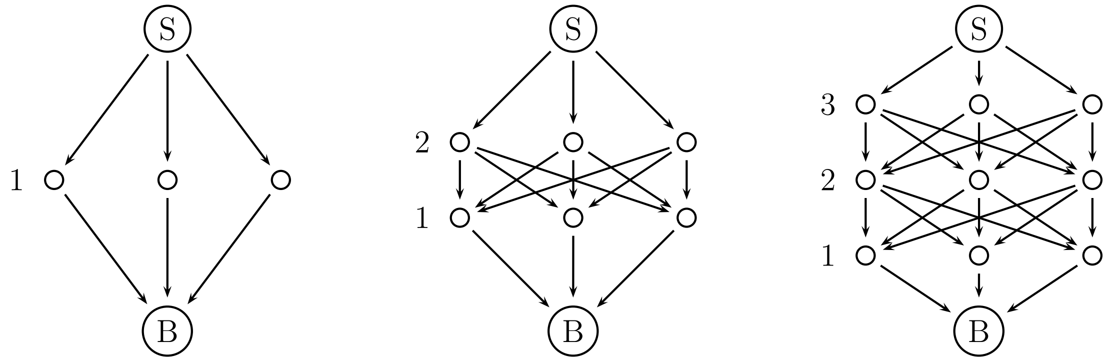

(a) A 1×3 network. (b) A 2×3 network. (c) A 3×3 network.

Figure 1: Example trading networks. The directed edges summarize the feasible transactions.

tially departing from their setting, we additionally assume that agents have private information, a common feature of many trading environments. As outlined in the introduction, we assume that intermediary traders face a private *liquidity constraint*, *budget constraint*, or *financial endowment* (we use the terms interchangeably). Thus, each trader has a private budget that he uses to finance his trading activity. When attempting to acquire the asset, each trader is constrained to bid less than his private budget. This budget has no direct impact on an agent's payoff; it only describes his feasible actions.

To model this constraint, we assume that endowments are independently and identically distributed according to the cumulative distribution function F. We assume that F admits a continuous density, f, with full support on the interval $[0, \bar{w}]$ where $0 < v \leq \bar{w}$. Though straightforward to relax, for analytical brevity we assume that the ratio $F(w)/f(w)$ is nondecreasing in w. Many common distributions, such as the uniform or the normal, satisfy this condition. As standard, this information structure is common knowledge among all agents in the economy, though a trader's realized budget (i.e., his type) is his private information. In the special case of an economy with one row of intermediary traders ($R = 1$) our model reduces to that of Che and Gale (1996).

3.1 Model Interpretation and Discussion

Having described our model, we wish to elaborate upon its underlying assumptions and associated interpretations. As explained below, our model captures the essential features

baseline experiment.

of intermediated exchange in a networked economy. It affords sufficient tractability for theoretical investigation and adequate simplicity for experimental scrutiny.

The Buyer and the Seller The focus of our analysis is intermediaries. Thus, we have de-emphasized the role of buyers and sellers. Indeed, there is but a sole seller and an equally solitary buyer. This restriction is entirely for expositional simplicity. The seller plays only an ancillary role in our model as a "source" of assets while the buyer serves as a corresponding "sink." Both agents may therefore serve as metaphors for larger, not modeled upstream and downstream markets. Hence, both are non-strategic actors, whose behavior is taken as given by traders.

The Network Structure A central feature of our setting is the network structure, which enjoys a characteristic "stacked" form. Choi et al. (2014) call such networks multipartite. Such trading networks capture several features of intermediated exchange at a micro level.

First, the number of rows—the parameter R—describes the prevailing depth of intermediation in the market. It counts the number of intermediaries involved in moving assets from sellers to buyers. Different markets will exhibit different values of this parameter. In Kranton and Minehart (2001), for example, $R = 0$ as buyers and sellers are linked directly. Often, however, R may assume a relatively large value. For example, Li and Schürhoff (2014) have recently documented the presence of long chains of successive intermediaries in over the counter (OTC) trades of municipal bonds. They note that approximately a quarter of transactions between buyers and sellers in this market involve multiple intermediaries. Describing these transactions, they write:

> [L]onger chains start with a dealer purchasing a bond from a customer, followed by one or several interdealer trades that move the bond from the head dealer to the tail dealer, and end with the tail dealer selling the bond to one or more customers. (Li and Schürhoff, 2014, p. 14)

This process closely aligns with our framework.

Second, the width of rows—the parameter N—captures the intensity of competition among intermediaries.[6] When competition is intense and N is large, many intermediaries are close substitutes in the provision of intermediary services. Hence, their rents get dissipated

[6]We assume a common value of N in each row. It is straightforward to allow this parameter to vary by row, as considered by Kotowski and Leister (2014). The qualitative characteristics of the economy's equilibrium will be unchanged.

accordingly. In contrast, when N is small, particular traders become more "essential" in linking dispersed buyers and sellers. Analogous to monopoly rents, "intermediation rents" rise accordingly. Many studies have investigated this interaction (Blume et al., 2009; Choi et al., 2014).

While stacked, multi-partite networks appear to form a particular class of markets; they also offer insights regarding more complex economies. Specifically, our networks capture well the *local interactions* among economic actors where inter-agent relationships are most likely to be salient. A stacked, multi-partite network thus has a natural interpretation as a small neighborhood within a larger (not modeled) economy. In this neighborhood intermediaries enjoy a dense set of relationships among each other with few ties further afield. Under this paradigm, the seller and the buyer can be interpreted as metaphors for larger upstream and downstream markets, respectively.

Private Budget Constraints Private information plays a central role in our model and in our experimental study. Introducing private information on top of an economic network, however, compounds the environment's complexity. We focus on the case of private budget constraints because they arise naturally in many markets and they capture, sometimes in reduced form, many types of market uncertainty. Beyond its literal interpretation as a hard financial constraint, a budget constraint may instead reflect an economic opportunity cost of funds.[7] While private budget constraints provide a flexible starting point for analysis, others have investigated alternative forms of private information in networked markets. For example, Kotowski and Leister (2014) consider a trading economy where intermediaries have private trading or inventory costs. Our focus is different, though complementary.

The Price-Setting Mechanism While it is possible to consider other price-setting protocols, we build our model around the first-price auction. It is a transparent and well-understood mechanism suitable for experimental investigation. By simplifying the price-setting process, we can focus on the network effects of trade with fewer confounds. Also in a bid to limit confounding effects, we have suppressed reserve prices (or ask prices) in the trading protocol. Their role has already been examined in the experiments of Gale and Kariv (2009), albeit in a perfect-information environment. As noted above, the experiments

[7]More abstractly, however, it may capture a trader's weak position in the network. A trader with weak relationships or ties to resale markets (modeled as links in the trading network) may be unwilling or unable to pledge a great sum to the particular market interaction. He might doubt the asset's quality or question potential buyers' interest. This case would be akin to a low realized "budget" that caps a trader's bid even though his hesitation is not directly linked to a low bank-account balance.

of Charness et al. (2007), Judd and Kearns (2008), and Choi et al. (2014) involve different trading protocols.

3.2 A Monotone, Pure-Strategy Equilibrium

We start by characterizing an equilibrium in monotone, pure strategies. Monotone, pure strategy equilibria are focal in auction-like settings. Che and Gale (1996), whose model our market generalizes, also focus on such equilibria. While the payment made by the buyer is fixed by assumption, the value of the asset to traders in rows $r \geq 2$ is determined endogenously. The value of the asset to a row-2 trader, for example, depends on the expected bids of row-1 traders. Likewise, the value of the asset to a row-3 trader depends directly on the expected bids of row-2 traders and indirectly on the bids of row-1 traders; and so on. The result of this inductive reasoning is an equilibrium featuring a recursive structure.

To specify the resulting equilibrium strategy, we first define $G(w) := F(w)^N$ as the cumulative distribution function (c.d.f) of the largest of N independent random variables, each with c.d.f. $F(w)$. As usual, $g := G'$ is the corresponding density.

Theorem 1. *Let $U_1(w) = \max_{b \in [0,w]} F(b)^{N-1}(v - b)$ and define*

$$b_1^*(w) = v - \frac{U_1(w)}{F(w)^{N-1}}.$$

For $r \geq 2$, define inductively the following expressions:

$$\nu_{r-1}^* = \int_0^{\bar{w}} b_{r-1}^*(w) g(w) \; dw \tag{1}$$

$$U_r(w) = \max_{b \in [0,w]} F(b)^{N-1} \left(\nu_{r-1}^* - b\right) \tag{2}$$

$$b_r^*(w) = \nu_{r-1}^* - \frac{U_r(w)}{F(w)^{N-1}} \tag{3}$$

The strategy profile where all traders in row r bid according to the strategy $b_r^*(w)$ is the unique, monotone, pure strategy equilibrium of the trading game.

The proof of Theorem 1 follows by induction from the analysis of Che and Gale (1996).

In the equilibrium of Theorem 1, a trader's bid is increasing in his endowment. Traders with low endowments bid all of their available resources in the contest. Their budget constraint binds. Those with higher endowments shade their bids relative to their available pool of funds. Such bidders are unconstrained. Below we characterize several additional features implied by this equilibrium. Before doing so, we provide an example paralleling our experimental investigation to follow.

Example 1. Suppose the buyer offers to pay $v = 100$ for the asset and traders' budgets are uniformly distributed on the interval [0,100]. Assume that $N = R = 3$, as in Figure 1(c). The equilibrium bidding strategies are illustrated in Figure 2. In closed form,[8] these bidding strategies are:

$$b_1^*(w) = \begin{cases} w & w < 67 \\ 100 - \frac{145000}{w^2} & 67 \leq w \end{cases}$$

$$b_2^*(w) = \begin{cases} w & w < 46 \\ 70 - \frac{50000}{w^2} & 46 \leq w \end{cases}$$

$$b_3^*(w) = \begin{cases} w & w < 39 \\ 59 - \frac{30000}{w^2} & 39 \leq w \end{cases}$$

The equilibrium strategy of agents in each row is strictly increasing in w and concave. As a group, equilibrium strategies are ordered with traders in row 1 submitting the highest bids.

Example 1 suggests several general implications of equilibrium behavior in our networked market. These conclusions touch on individual-level strategies, price dynamics, and welfare implications. We summarize these features in three corollaries, which will inform our experimental investigation.

Focusing on bidder behavior, we observe two striking features in the strategies identified in Example 1. First, bidders' strategies are increasing and concave. For bidders in each row, there is a critical endowment level, which we denote by w_r^*, that defines a boundary between two types of bids. For budgets below the critical value, the equilibrium strategy coincides with a simple heuristic—"bid everything." For budget levels above the critical value, equilibrium bidding demands a more sophisticated approach whereby a bidder commits only a fraction of his endowment to the trading game.

Second, bidders' strategies are ordered uniformly. Conditional on endowment, traders in row 1 bid more aggressively than row 2 traders. Row 2 traders bid more aggressively than row 3 traders. This ordering reflects the changing value of the asset as it proceeds through the network. In the example, Row 1 traders can resell the asset to the buyer for a price

[8] For expositional ease, we round several (otherwise exact) values.

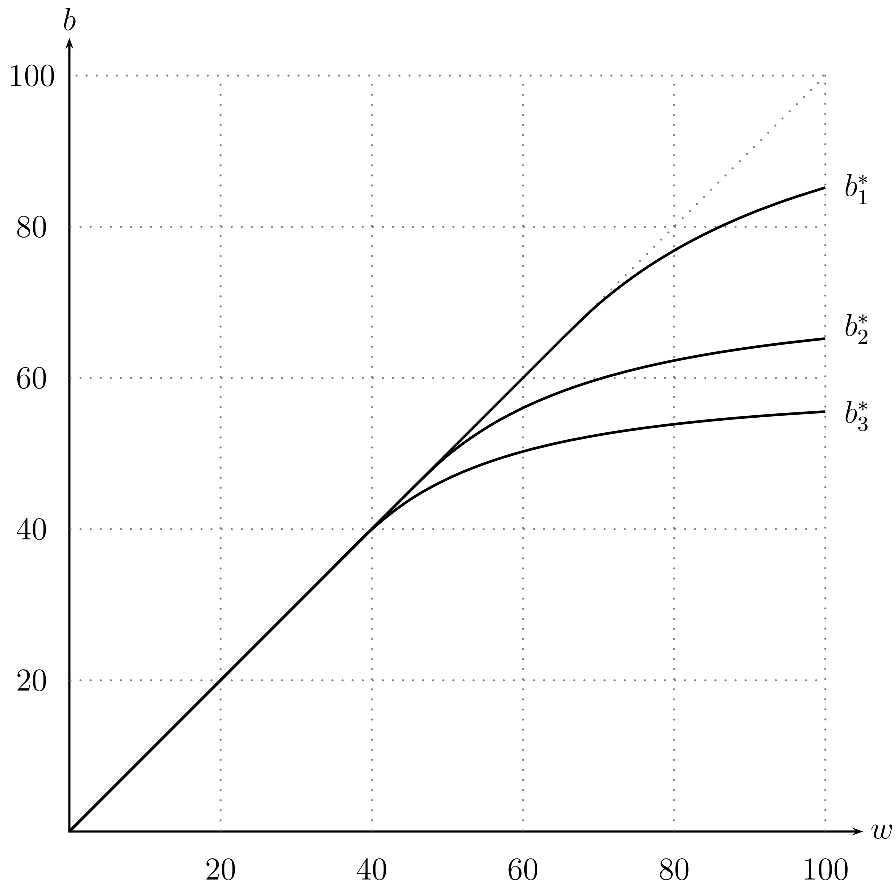

Figure 2: Equilibrium bidding strategies in Example 1.

of 100. A trader in row 2, however, can only resell the asset to a trader in row 1. Since row 1 traders strategically shade their bids relative to their budget and realized budgets are almost surely less than 100, the expected resale price is only approximately 70. Thus, a row 2 trader should shade his bid more aggressively than a row 1 trader. By induction, the same reasoning applies to traders in subsequent rows moving away from the buyer. These two features are in fact general properties of this market's equilibrium.

Corollary 1 (Bidding)**.** *Consider the equilibrium characterized by Theorem 1.*

1. *There exists a critical value w_r^* such that for all $w \leq w_r^*$, $b_r^*(w) = w$. For all $w > w_r^*$, $b_r^*(w)$ is strictly increasing, concave, and strictly less than w.*

2. *Conditional on his endowment level, traders further from the buyer bid less, i.e. $b_r^*(w) \geq b_{r+1}^*(w)$ for all w.*

Building on Corollary 1, we can derive equally clueful conclusions concerning equilibrium price dynamics. Given our trading protocol, this value equals the highest bid submitted in each respective auction. Since $b_r^*(w)$ is strictly increasing, the highest bid is always submitted by the trader with the highest realized endowment. Thus, and following directly from Theorem 1, realized prices within a trading network are independent and $\nu_r^* = \int b_r^*(w)g(w)dw$ is the expected price paid by traders in row r. This value evolves predictably through the network.

Corollary 2 (Prices)**.** *Consider the equilibrium characterized by Theorem 1.*

1. *Expected prices are increasing as the asset approaches the buyer, $\nu_{r+1}^* \leq \nu_r^*$.*

2. *The difference in expected prices is increasing as the asset approaches the buyer, $\nu_r^* - \nu_{r+1}^* \leq \nu_{r-1}^* - \nu_r^*$.*

While the first point of Corollary 2 is intuitive, the second is surprising. As the buyer's offer of v acts as a hard bound on expected prices and these prices are increasing, we might expect that they "converge" to the buyer's value with later traders capturing smaller-and-smaller shares of the total surplus. Instead, prices evolve in the opposite manner. Among all intermediaries, traders closer to the buyer capture the greatest share of the expected total surplus. (Of course, part of the surplus goes to the seller as well.)

As prices evolve systematically through the trading network, it is natural that traders' payoffs behave equally predictably. From the bidding strategy defined in Theorem 1, we can conclude that the equilibrium expected payoff of a trader with an endowment of w equals $U_r(w)$, as defined in (2). Indeed, if we rearrange (3), we see that

$$b_r^*(w) = \nu_{r-1}^* - \frac{U_r(w)}{F(w)^{N-1}} \implies U_r(w) = F(w)^{N-1}(\nu_{r-1}^* - b_r^*(w)).$$

Conditional on acquiring the asset, a trader earns an expected return of $\nu_{r-1}^* - b_r^*(w)$, the difference between the expected resale price and his bid. The probability with which he acquires the asset equals $F(w)^{N-1}$, which is simply the probability that all others in his row have an endowment less than w and bid less than $b_r^*(w)$. With this description of payoffs, we can conclude the following.

Corollary 3 (Payoffs)**.** *Consider the equilibrium characterized by Theorem 1. Define $U_r(w)$ as in* (2) *and define w_r^* as in Corollary 1.*

1. *$U_r(w)$, is strictly increasing in w for all $w \leq w_r^*$. For all $w > w_r^*$, $U_r(w)$ is constant and equal to $U_r(w) = F(w_r^*)^{N-1}(\nu_{r-1}^* - w_r^*)$.*
2. *Expected trading profits are increasing as the distance to the buyer decreases, $U_{r+1}(w) \leq U_r(w)$.*

Corollary 3 shows that traders' endowment-contingent expected payoffs differ systematically through the network. Greater profits are earned by traders closer to the buyer. As all traders are ex ante identical, except for their position in the network, this implies that unconditional expected payoffs are also greatest for traders closer to the buyer. Despite the analogous network topology, these observations contrasts with the conclusions of Gale and Kariv (2009), who argue that intermediation rents are competed aways. In our setting, the presence of private information ensures that these rents remain positive.

3.3 Mixed and Other Equilibria

Though the analysis above, and our experimental investigation below, focus on the monotone, pure strategy equilibrium, there are in fact many *mixed* strategy and non-monotone equilibria that share an equivalence with the equilibrium identified in Theorem 1. Briefly, many other strategy profiles generate the same distribution of bids and, by straightforward reasoning, can also be shown to be equilibrium strategies.

To appreciate this observation, recall that Corollary 3 concluded that the equilibrium payoff of a trader with budget $w > w_r^*$ is

$$U_r(w) = F(w_r^*)^{N-1}(\nu_{r-1}^* - w_r^*),$$

which is a constant value independent of his bid (above w_r^*). Therefore, given others' strategies, a trader in row r is indifferent among all bids in the range of $b_r^*(w)$ above w_r^*. Conditional on having a sufficient budget, an intermediary can plan to randomize his bid in this range. Provided that across bidder types the randomization replicates the equilibrium bid distribution induced by the strategy in Theorem 1, another equilibrium will obtain. We elaborate on this fact in Online Appendix A where we describe such mixed strategy equilibria in greater detail as equilibria in distributional strategies (Milgrom and Weber, 1985).

Remark 1. The conclusions of Corollaries 2 and 3 continue to apply whenever the distribution of bids replicates the distribution of bids from the monotone, pure-strategy equilibrium of Theorem 1. Hence, those conclusions are robust even if the underlying equilibrium is asymmetric, non-monotone, or in mixed strategies. This ensures that many of the empirical predictions identified below are independent of the selected equilibrium.

4 Model Predictions and Experimental Procedures

To organize our experimental investigation, we will first translate our formal results into a collection of empirical predictions that we will take to the laboratory. Thereafter, we explain our experimental design.

First, Theorem 1 characterized a pure strategy equilibrium in our trading model. Together with Corollary 1, this equilibrium exhibits several characteristic features. As monotone equilibria are focal in many market games, we can translate those conclusions into the following qualitative prediction concerning anticipated bidding behavior.

Prediction 1 (Monotone Bidding)**.**

1. Within a row, traders with greater endowments will place greater bids; however, traders with high endowments will bid less than their budget.

2. Conditional on endowment level, traders closer to the buyer will exhibit more aggressive bidding behavior.

Table 1: Summary of experimental sessions and treatments

Network Type	Number of Sessions	Total Subjects	Decision Observations	Average Earnings ($)*
1×3	1	30	1,500	21.04
2×3	2	72	3,600	19.51
3×3	3	90	4,500	16.03

* Includes the $10 participation fee and earnings from the experiment.

Second, Corollary 2 describes the evolution of prices in the economy. As they are functions of agents' endowment levels, prices should be independent, but they should exhibit an upward trend on average. The highest average prices should be paid by traders closest to the buyer and the lowest average prices should be paid by traders closest to the seller.

Prediction 2 (Prices)**.** The prices paid by intermediary traders in different rows exhibit an upward trend as the asset approaches the buyer. The gap between purchase and resale prices is increasing as the asset approaches the buyer.

Our final prediction recasts Corollary 3, establishing comparative statics on traders' welfare and payoffs within and across rows.

Prediction 3 (Payoffs)**.**

1. Within a row, average intermediary profits are increasing at low endowment levels and constant for high endowment levels.

2. Conditional on the endowment level, traders closer to the buyer earn higher average payoffs.

While Prediction 1 directly builds on the monotone, pure-strategy equilibrium, Predictions 2 and 3 apply to the large family of (possibly asymmetric) equilibria, as explained in Remark 1 above.

4.1 Experimental Design and Procedures

In our experiment, we consider a family of networks to test our main predictions. Table 1 provides a summary. First, the 1×3 network allows us to examine the *within* row elements of Predictions 1.1 and 3.1. This network, which corresponds to a test of Che and Gale's (1996) model of a common-value, first-price auction with private budget constraints, allows us to identify any dependencies in agents' bids and payoffs on endowment levels with minimal confounds. The 1×3 network is insufficient to probe the price and welfare dynamics *across* rows. Thus, we also consider the 2×3 network and the 3×3 networks, having two and three layers of intermediaries, respectively, between the buyer and the seller. These more complex networks also provide a forum to ascertain the robustness of our within-row conclusions. The monotonicity, concavity, and welfare conclusions apply to these cases as well. To ensure comparability, we maintain the same level of competition across treatments. Thus, each row always has three intermediary traders.

Six experimental sessions were conducted at the Experimental Social Science Laboratory (Xlab) at the University of California, Berkeley. We employed the experimental computer program developed by Gale and Kariv (2009) and our experimental procedures follow their's closely. Subjects were recruited from the undergraduate and graduate student bodies at the university. No subject participated in more than one session. At the start of each session, subjects were given ample time to silently read the experiment's instructions. We reproduce a sample of these instruction in Online Appendix B. The instructions were then read aloud by an administrator. Subjects' questions, if any, were publicly addressed. Each session lasted roughly one hour and thirty minutes. A \$15 participation fee plus earnings from the experiment were paid in private at the end of each session.[9] Earnings in the experiment ranged from zero to \$72.90.

Each of six sessions consisted of 50 independent rounds of trade. Subjects assumed the role of intermediaries and each round mimicked our model's operation. The buyer and the seller were modeled by the computer. The network structure was held fixed throughout a session and was clearly explained to the subjects. Each subject was randomly assigned by the computer to one row of the network. These assignments were privately observed by the subjects, and were held constant throughout the experiment's duration. At the beginning of each trading round, the computer randomly formed networks by assigning subjects to their respective row in a network. Subjects faced the same probability of being placed to any network and to any node position within their preassigned row. Thus, subjects were unaware of the identities of others in their network in each round.

Upon observing their position in the network and immediately prior to each round of play, each subject was informed of his trading endowment for that round. This endowment was randomly and uniformly drawn between 0 and 100 tokens (including decimal values) and

[9]The participation fee comprised of \$5 paid for arriving plus \$10 paid for completing the experiment.

not carried over between rounds. Endowments were independently drawn across subjects. Each subject was then asked to choose a price (bid) at which they were willing to buy the asset. Each subject was allowed to choose any number (including decimals) between 0 and their endowment as their bid. Once all bids had been submitted, trades were executed by the computer as follows. The subject connected to the (computer) seller who submitted the highest bid was awarded the asset and automatically resold the asset to the trader submitting the highest bid in the row below. This process continued until the holder of the asset sold it to the (computer) buyer, who always paid 100 tokens. This trading procedure was explained to the subjects prior to the experiment.

After all trades in a round were executed, subjects learned the outcome. This information was conveyed via the experimental platform's interface, which we now describe. Figure 3 presents a typical dialog window from a session with a 3×3 network. The interface for the other network structures was similar. The subject's position in their network are displayed in the large square window at the left of the screen. Upon completion of a given round, this window also displayed that round's price and trading information. The "View Results" button in the screen's top-right corner allowed the subject to view the price and trading history from each previous trading round they participated in. The subject entered his or her bid in the Bid field at the right of the screen. The subject's endowment was displayed immediately above the Bid field. After each round of bidding, each subject was informed of others' bids and endowments. They were also shown the sequence of executed trades and the associated prices, as illustrated in Figure 3. Provision of this information ensured that subjects trusted and understood the trading protocol and pricing system. After each subject clicked "OK," the next trading round began.

Earnings in the experiment were determined as follows. At the end of the session, the computer randomly selected one trading round, where each round had an equal probability of being chosen. The subjects were then paid based on that round's earnings. If a subject did not trade in the selected round, their trading profits were zero and they only received a \$10 participation fee for completing the experiment. If a subject did trade in the selected round, they received a \$10 participation fee *plus* their earnings from the selected round. These earnings were computed to be

$$\text{Earnings} = \max\left\{\underbrace{\text{Resale Price} - \text{Own Bid}}_{\text{Trading Profits}}, 0\right\}. \tag{4}$$

As a subject's trading profits can in principle be negative, calculating earnings with formula

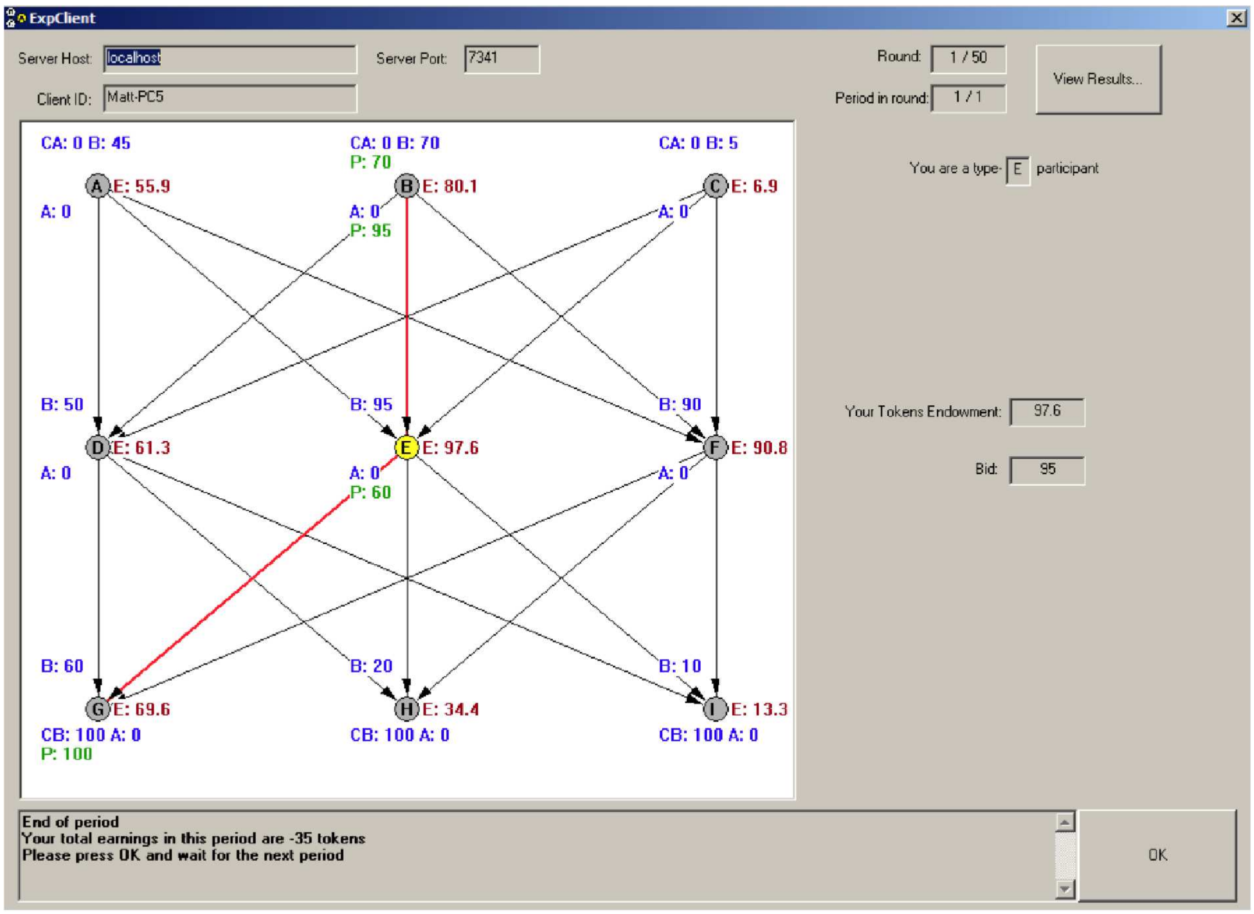

Figure 3: Sample interface in the 3 × 3 network treatment.

(4) ensured that subjects did not incur a loss in the experiment by introducing a limited liability constraint.[10] As earnings formula (4) differs slightly from our baseline model developed in Section 3, we precisely characterize its implications in the subsection below. We confirm that it has no effects on the qualitative features of our economy's equilibrium and its quantitative implications are mild, predictable, and can be easily controlled for, when required. Subjects were informed of the earnings formula as part of the experiment's instructions. To compute final payments, tokens were converted into dollars at a 1-for-1 exchange rate. As noted in Table 1, average earnings are greatest in the 1×3 network and lowest in the 3×3 network. This difference is consistent with our theoretical model since row 2 and row 3 traders have lower predicted payoffs than row 1 traders (Prediction 3).

4.1.1 Limited Liability and Experiment Earnings

A practical challenge that our experimental investigation must address is that an intermediary trader may incur an ex post loss. With the exception of traders in row 1, it is possible that a trader pays more for the asset than he receives from a downstream counter-party. Similar concerns commonly feature in studies of common-value auctions where bidders may be vulnerable to the winner's curse. In those experiments, subjects who overestimate the item's value overbid precipitating negative earnings or bankruptcy. Recognizing the practical complications losses impose within an economic experiment, numerous methods have been employed in the literature to resolve these concerns, sometimes with mixed effects. Kagel (1995) provides a survey.

Following Gale and Kariv (2009), we address bankruptcy risk and subject losses *directly* by incorporating a limited-liability constraint into our earnings formula, as described above.[11] Unlike Gale and Kariv (2009), we account for the additional constraint by developing a parallel theoretical analysis investigating its implications. Importantly, as explained below, the limited liability correction leaves unaffected the qualitative conclusions identified in Corollaries 1–3 and has only mild and predictable quantitative implications for the equilibrium identified in Theorem 1. Predictions 1–3 continue to apply.

[10] Our experimental protocol and laboratory rules necessitated this requirement. An alternative experimental design could have endowed bidders with a "buffer" budget of tokens from which trading losses could be deducted. The buffer could have been chosen to be sufficiently large to ensure that ex post losses are extremely unlikely. We decided against such a design due to its added cost and complexity. Notably, a subject would need to keep track of both their buffer budget and their private budget constraint in a particular round. We believe the present design is more closely aligned with our original model.

[11] Gale and Kariv (2009) claim that payoff formula (4) be interpreted as a compensation scheme for a professional trader who enjoys bonuses when profits are positive, but earns constant wages otherwise.

When earnings are computed according to formula (4), incentives are impacted. Since traders are protected from a loss, intuition suggests that bidding in equilibrium ought to be more aggressive than otherwise. Theoretical analysis confirms that this is case. In Online Appendix C we provide an equilibrium characterization in our economy, analogous to Theorem 1, when (4) defines traders' payoffs. Though considerably less tractable than our baseline model from Section 3, equilibrium play can again be defined inductively. We can compute (numerically) the equilibrium strategy given our experimental parameterization to gauge the impact of employing earnings formula (4) vis-à-vis our baseline model.

Carrying out the described analysis yields bidding strategies reminiscent of those from our original model. In Figure 4 we illustrate the computed strategies as dashed curves. These strategies are labeled as $\tilde{b}_r(\cdot)$ to distinguish them from the equilibrium strategies of our baseline model, which are labeled as $b_r^*(\cdot)$. The bidding strategy of row 1 traders is unaffected by the introduction of limited liability since the buyer always offers more for the asset than the highest possible bid. Thus, $\tilde{b}_1(w) = b_1^*(w)$. As expected, traders in rows 2 and 3 adopt more aggressive postures at each endowment level than originally. They are willing to bid more since the down-side risk associated with overpaying is mitigated by formula (4). Consequently, $\tilde{b}_2(w) \geq b_2^*(w)$ and $\tilde{b}_3(w) \geq b_3^*(w)$ as illustrated in Figure 4. Like in the baseline case, for each row r we define $\tilde{w}_r$ as the critical endowment level such that $\tilde{b}_r(w) = w$ for all $w < \tilde{w}_r$ and $\tilde{b}_r(w) < w$ for all $w > \tilde{w}_r$.[12]

Importantly, the implications of these adjustments are mild. All qualitative conclusions outlined in Corollaries 1–3 continue to obtain. The continued applicability of Corollary 1 is apparent in Figure 4. Within our parameterization, Corollaries 2–3 can be easily verified, either analytically or numerically. Thus, Predictions 1–3 continue to be valid. Maintaining applicability of our original conclusions, (4) provides a practical, minimally invasive compromise in implementing our experimental market.

5 Data Analysis and Results

As a starting point for our data analysis, we examine bidder behavior. In Figure 5 we plot LOESS curves (and 95% confidence bounds) summarizing average bids submitted by traders in each row as a function of their endowment level.[13] We group our estimates by

[12] In Example 1, the critical values (w_r^*) are approximately 39, 46, and 67. With limited liability in the payoff formula, the critical values ($\tilde{w}_r$) become 43.5, 50.5, and 67, respectively.

[13] LOESS is a standard non-parametric regression procedure, which fits a local polynomial regression to (in our case) bivariate data. See Fox (2002) for a practical discussion of this and related methods.

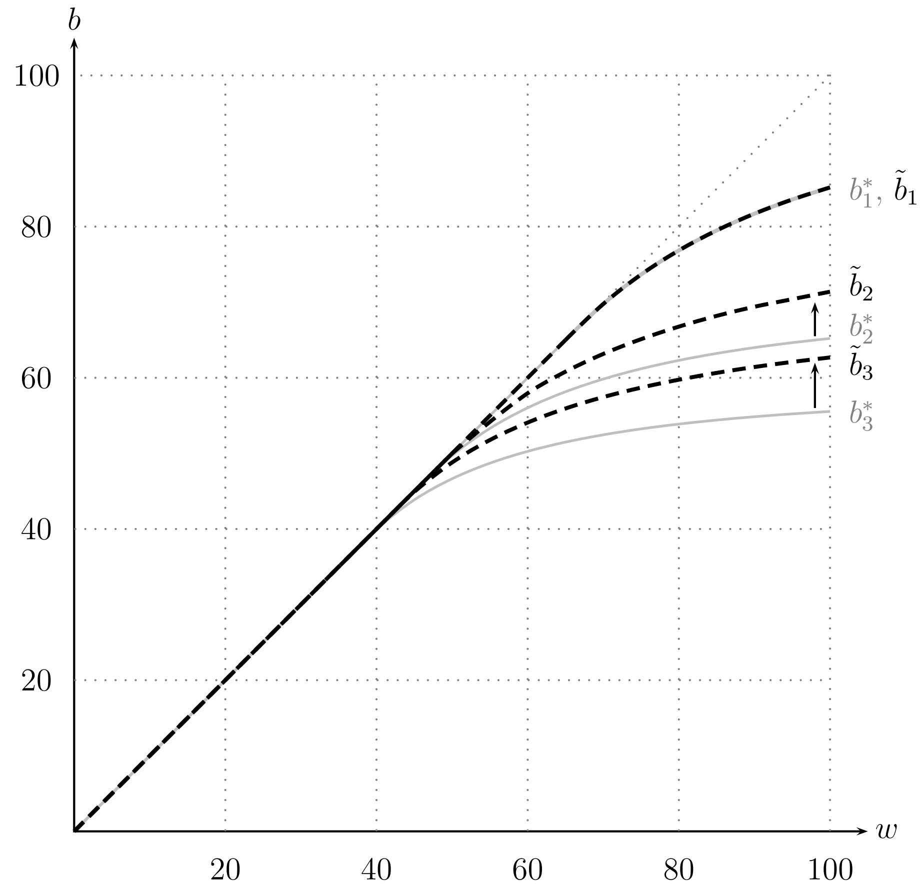

Figure 4: Monotone, pure-strategy equilibrium strategies without limited liability (b_r^*) and with limited liability ($\tilde{b}_r$) in the economy of Example 1. Limited liability increases the bids of traders in rows 2 and 3. The strategy of row 1 traders is unchanged.

network type, but we pool across sessions and bidders.[14] In each successive panel, we consider increasingly complex networks, with the 3×3 network presented in panel (c).

Two features of the fitted curves stand out. First, the estimates exhibit the anticipated monotonicity and concavity predicted by our theoretical analysis. Across networks, traders bid their entire budget when their realized endowments are low. At higher endowment levels, traders do not commit their entire budget to the auction. Instead, they shade their bid strategically resulting in a concave bidding function. Second, the fitted curves conform to the predicted ordering. Row 1 traders uniformly outbid their row 2 and row 3 counterparts. Moreover, in the 3×3 treatment, row 2 traders outbid traders in row 3. The associated confidence bounds overlap only at extreme values where, due to fewer observations, estimates are necessarily less precise.

Result 1 (Bidding)**.** Across networks, traders' bids are increasing in endowment, even when endowments are not a binding constraint. Traders with low endowments bid their entire endowment to acquire the asset. Traders with high endowments strategically reduce their bid relative to the their budget. Across networks, intermediaries closer to the final buyer adopt uniformly more aggressive bidding strategies.

As for price dynamics, in Table 2 we report descriptive statistics concerning the prices paid by traders to acquire the asset. The table summarizes these values conditional on the network structure and on the row, though the values are pooled across sessions and trading rounds. For example, a trader in a 1×3 network who successfully acquired the asset from the seller, paid 70.71 tokens on average. The median price paid was higher, 75 tokens, and the standard deviation was 17.16 tokens. Similar statistics are reported for other networks on a row-by-row basis.

A first-order feature apparent from Table 2 is the trend in average prices as the asset approaches the buyer. In a 2×3 network, the price paid by a row 1 trader to acquire the asset (from a row 2 trader) is on average 10.5 tokens greater than the price paid by a row 2 trader to acquire the asset (from the seller). In a 3×3 network, a row 1 trader paid on average 7.18 tokens more than a row 2 trader. A row 2 trader paid on average 2.57 tokens more than a row 3 trader. Similar differences exist among median prices. These observed differences are (statistically) distinct from zero. For example, t-tests and sign tests reject the null hypothesis of equality with p-values far below 0.001.

[14] In Online Appendix D we estimate a parametric model controlling for subject-level effects. The conclusions are entirely analogous to those presented here. Pooling by session also yields similar conclusions.

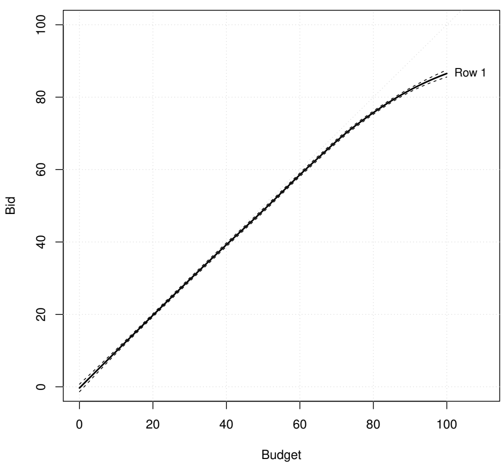

(a) Estimated bidding strategy in the 1×3 network.

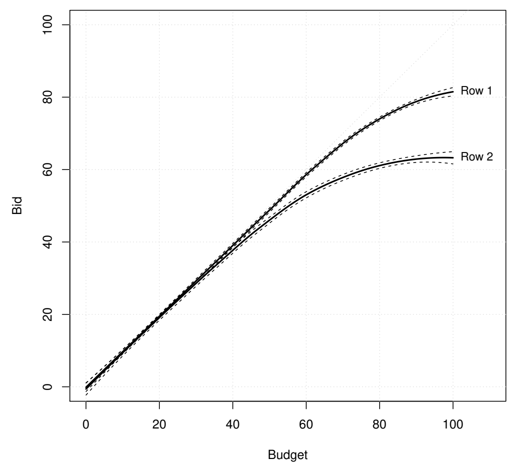

(b) Estimated bidding strategy in the 2×3 network.

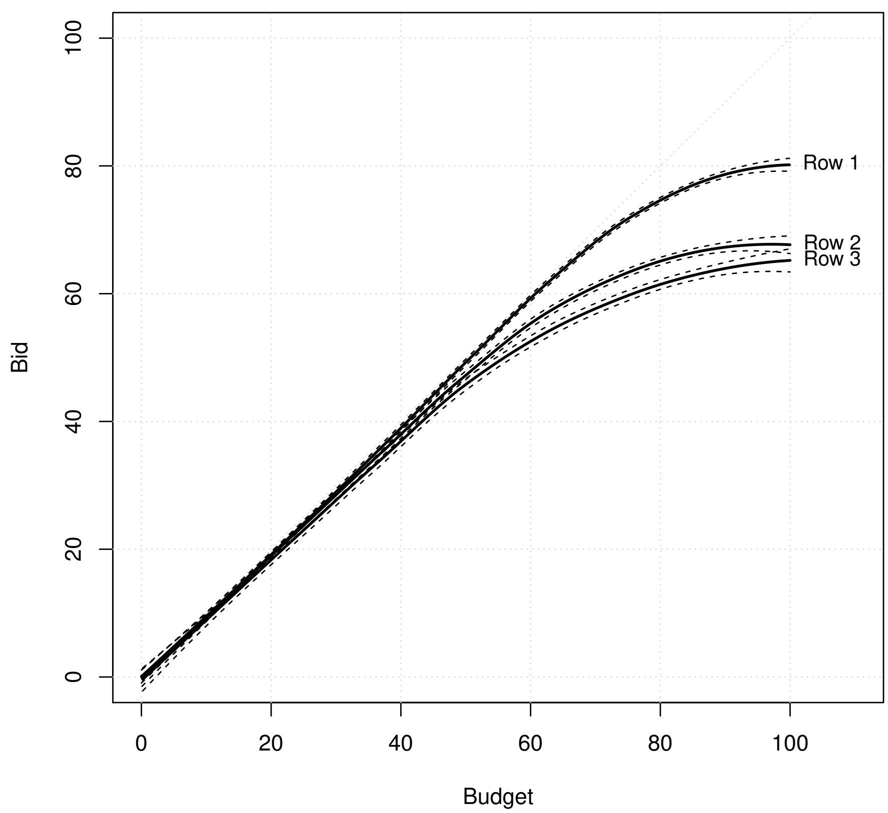

(c) Estimated bidding strategy in the 3×3 network.

Figure 5: Estimated bidding strategies. LOESS curves with 95% confidence bounds.

Table 2: Summary statistics of trading prices.

Network	1×3	2×3		3×3		
Row	1	1	2	1	2	3
Mean	70.71	68.91	58.41	68.51	61.33	58.76
Median	75.00	71.30	60.00	72.50	62.10	60.10
St. Dev.	17.16	17.22	14.82	16.02	13.93	15.62
Obs.	500	600	600	500	500	500

As noted in Corollary 2, our model also provides a prediction concerning second-order price dynamics. The gap in realized prices between adjacent rows is predicted to increase as the asset approaches the buyer. Above we noted why this prediction may challenge initial intuition given the mechanical "cap" on resale prices provided by the buyer's constant offer price. To test this prediction, we again consider the 3×3 network. As noted above, in that economy we observed notable differences in the average realized prices between adjacent rows. Similarly, the median difference in realized prices was 9.1 between rows 1 and 2 and and 2.4 between rows 2 and 3.[15] Again this difference is statistically distinct from zero. For example, a sign test rejects equality ($p = 0.043$, two-sided).

Result 2 (Prices)**.** Average prices are increasing as the asset approaches the buyer. Moreover, the gap between purchase and resale prices is increasing.

Our analysis has also provided predictions concerning the evolution of profits across rows, conditional on an endowment level, and across endowment levels, conditional on row. To illustrate both comparative statics, in Figure 6 we plot bidders' payoffs as a function of realized endowment levels stratifying the sample by row, but pooling across subjects. To summarize general trends, our plots again consist of LOESS curves, which provide an estimate of the average payoffs conditional on endowment level. Payoffs are computed accounting for the limited liability constraint as given by formula 4 and explained in Section 4. This correction is necessary to ensure that we correctly capture the incentives faced by the subjects in the experiment. Panel (a) provides a succinct summary of general trends by combining data from all network types and sessions. Panels (b)–(c) disaggregate the data by network type.

Agreeing with Prediction 3, traders closer to the buyer enjoy higher average payoffs conditional on realized budget. Empirically, the sole exception to this general pattern occurs

[15]The median difference in realized prices differs from the difference in median realized prices, which can be computed from Table 2.

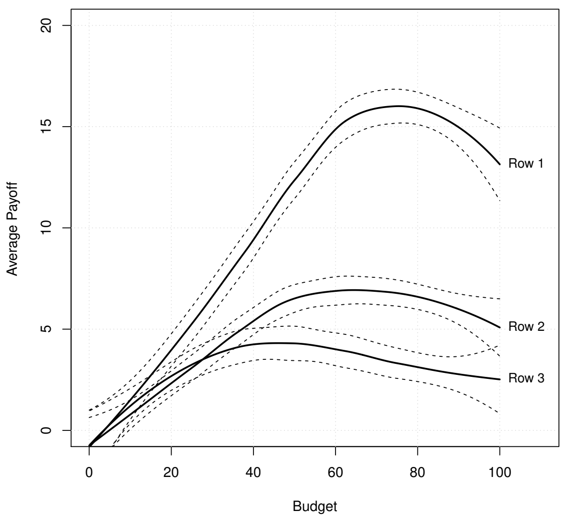

(a) All networks.

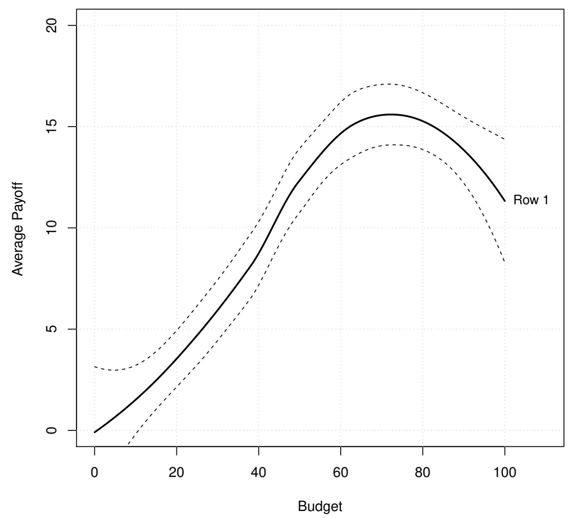

(b) 1×3 Network.

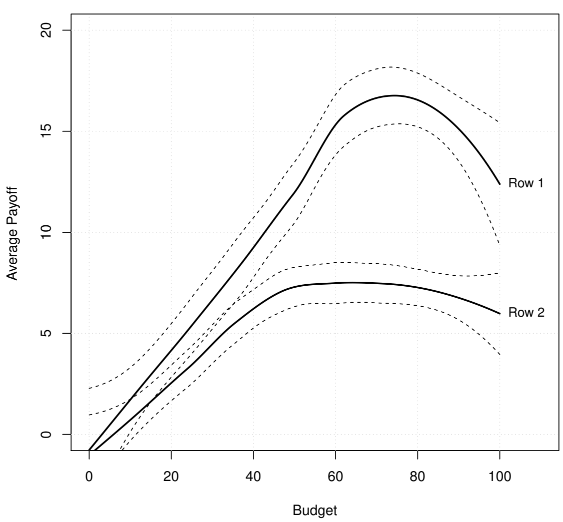

(c) 2×3 Network.

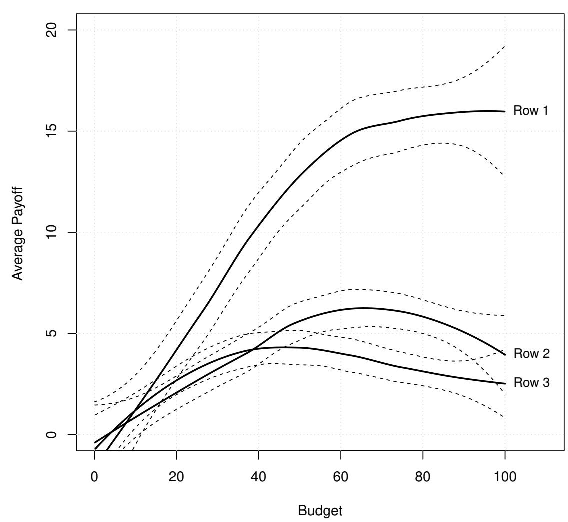

(d) 3×3 Network.

Figure 6: Average earnings as a function of endowment. LOESS curves with 95% confidence bounds.

at low endowment levels. Traders with low endowments always have a small chance of successfully acquiring the asset; hence, payoffs are (unsurprisingly) uniformly low. At high endowment levels, however, the ordering of traders' payoffs is evident and robust. It appears in the aggregate data and on a network-by-network basis.

Though the evolution of payoffs across rows provides a definite comparison, the evolution of payoffs within a row is more subtle. Recalling Corollary 3 and Prediction 3, we anticipate that average payoffs will be increasing when budgets are low and constant when budgets are high. The former effect is apparent all panels of Figure 6. The latter effect, however, appears to not be a feature of the data. Curiously, average payoffs exhibit a mild downward trend at high budget level. This effect is evident in panel (a), though it is a persistent feature at a finer scale, with the exception of row 1 bidders in the 3×3 network illustrated in panel (d).

To evaluate the significance of and to quantify this apparent downward trend, we first draw on our theoretical model to aid in sample selection. Restricting the sample to budget realizations above $\tilde{w}_r$, we regress trader's earnings on their realized budget. (Recall that $\tilde{w}_r$ denotes the cut-off value after which traders bid less than their endowment in equilibrium and expected payoffs are in theory constant.) By Corollary 3 and Prediction 3, the resulting regression estimates should be zero. Tables 3 and 4 report the outcomes of these calculation controlling for subject fixed effects. For comparison, Table 3 corresponds to Figure 6, panel (a), and Table 4 corresponds to Figure 6, panels (b)–(d).

In Table 3 we identify a mild, though statistically significant, negative correlation between realized budgets and payoffs. Thus, beyond the $\tilde{w}_r$ threshold, increasing a trader's endowment leads to a decline in average earnings. For example, increasing the endowment of a row 1 trader from 67 to 100, thereby removing all constraints on his bid, reduces his average payoff by about 3 tokens. Similar declines affect row 2 and row 3 traders. This effect appears despite the limited liability constraint present in our experiment, which precludes trading losses from pulling-down realized payoffs further.[16] Table 4 repeats the estimates of Table 3, but on more disaggregated samples. Across all specifications, the consistently negative point estimates of the budget coefficient corroborate the presence of a downward trend, though some estimates lose statistical significance.

We interpret this mild downward trend as a strategic analogue of the well-known winner's curse. Both effects besiege bidders with unusually "high" private information and complicate their strategic calculus. In our case, the blessing of a high budget comes with the curse of

[16] Neglecting this constraint maintains the ordering of agent's payoffs but significantly reinforces the downward trend in the payoffs of row 2 and row 3 traders at high endowments.

Table 3: Earnings trends among high-budget traders. OLS regression with subject fixed effects. Pooled sample across all networks.

	Row 1	Row 2	Row 3
Constant	22.834**	9.249**	10.834**
	(3.771)	(2.517)	(2.589)
Budget	-0.099**	-0.042*	-0.039*
	(0.035)	(0.019)	(0.016)
Obs.	1,571	1,552	834
Cut-off ($\tilde{w}_r$)	67	50.5	43.5

Robust standard errors in parenthesis. * $p < 0.05$. ** $p < 0.01$.

Table 4: Payoff trends among high-budget traders. OLS regression with subject fixed effects.

Network	1×3	2×3		3×3		
Row	1	1	2	1	2	3
Constant	24.996**	25.301**	7.932**	17.077**	10.244**	10.834**
	(6.149)	(5.570)	(3.244)	(5.794)	(2.901)	2.589
Budget	-0.120*	-0.137*	-0.030	-0.031	-0.055*	-0.039*
	(0.058)	(0.059)	(0.027)	(0.062)	(0.025)	(0.016)
Obs.	489	594	831	488	721	834
Cut-off ($\tilde{w}_r$)	67	67	50.5	67	50.5	43.5

Robust standard errors in parenthesis. * $p < 0.05$. ** $p < 0.01$.

more opportunities for error. An intermediary with a low endowment faces a relatively straightforward problem in formulating his bid. With very high probability the asset's resale value will substantially exceed his endowment. Bidding up to his budget constraint provides a simple, focal heuristic strategy guaranteeing considerable upside in the (admittedly, unlikely) event that he is the winner. Limited liability shields him from losses while the budget constraint provides protection from overbidding.

The problem faced by a high-budget bidder is more complex. Bidders with high budgets have sufficient funds to compete for the asset. Hence, they must be more careful in balancing the probability of winning the auction and the surplus derived from resale. This introduces additional channels for error or misperception to enter into the bid-formulation process. A high budget provides no mechanical backstop preventing overbidding thereby allowing these miscalculations and temptations to filter into the trading game.

Result 3 (Trader Profits)**.** On average, intermediaries closer to the buyer earn higher trading profits. Average profits are increasing in endowment when endowments are low. At high endowment levels, further increases in a trader's budget are associated with a mild, though statistically significant, decline in a trader's earnings.

6 Conclusion

We study a decentralized market with liquidity and financial risk. A multipartite network defines the set of feasible trades. The number of rows within the network, each comprised of a set of competing traders, defines the extent of intermediation required to successfully transfer an asset to a buyer. We identify a monotone, pure-strategy equilibrium with intuitive features. Traders account for others limited financial capacity by strategically shading their bids, aiming to acquire the asset at a bargain. Limited financial capacity in resale markets, particularly amongst traders in distant parts of the network, compounds the complexity and strategic uncertainty. As a result, (average) prices rise with successive transactions and intermediaries positioned closer to the buyer enjoy greater expected profits.

We empirically investigate our theory's key predictions in the laboratory. We observe pricing and payoff dynamics that are broadly consistent with the theory. Strategic uncertainty builds through the network in a predictable manner, with traders closer to the buyer realizing greater average payoffs. Aggregate bidding behavior aligns closely with the monotone equilibrium. We also observe a mild negative correlation between expected profits and budgets across high budget realizations. This holds within each row and is persistent across network treatments. Liquidity-rich traders tend to overbid. This tendency is mitigated by budget constraints and underscore their disciplinary role in markets, preventing costly "trembles" or "errors."

While the implication of private liquidity risk has been well studied in centralized exchange,[17] the pricing and distributional implications of liquidity risk in decentralized markets is much less understood. Our analysis provides clear predictions, substantiated by experimental data, concerning the compounding of risk in a decentralized market. We also describe the strategies adopted by market participants in response to its presence. Networks provide a systematic yet flexible framework for modeling decentralized market structures. More work is needed to further examine the interplay of liquidity and financing risk and a market's network structure. Analyzing specific policy interventions, such as liquidity injections that target trading positions within the market, offers a promising avenue for future research.

7 Proofs

Proof of Theorem 1. The proof is by induction. The base case for b_1^* follows immediately from Che and Gale (1996). For $r > 1$, the expected resale value of the asset is ν_{r-1}^*. Applying the argument from Che and Gale (1996) to a first-price auction with a common value of ν_{r-1}^* confirms that $b_r^*(w)$ is an equilibrium strategy for intermediaries in row r. Their argument applies as $b_r^*(w)$ is independent of the asset's price history. □

Proof of Corollary 1.

1. Since $F(w)/f(w)$ is increasing, $F(w)^{N-1}\left(\nu_{r-1}^* - w\right)$ is concave and attains a maximum at a unique value, which we call w_r^*. Thus, when $w < w_r^*$, $U_r(w) = F(w)^{N-1}(\nu_{r-1}^* - w)$ and (3) reduces to $b_r^*(w) = w$. When $w > w_r^*$, $U_r(w) = F(w_r^*)^{N-1}(\nu_{r-1}^* - w_r^*)$. Hence,

$$b_r^*(w) = \nu_{r-1}^* - \frac{F(w_r^*)^{N-1}(\nu_{r-1}^* - w_r^*)}{F(w)^{N-1}}.$$

From the first-order conditions that define w_r^*, $\frac{F(w_r^*)}{f(w_r^*)} = (N-1)(\nu_{r-1}^* - w_r^*)$. Thus, ${b_r^*}'(w) = F(w_r^*)^{N-1}(\nu_{r-1}^* - w_r^*)(N-1)f(w)F(w)^{-N} = \frac{F(w_r^*)^N}{f(w_r^*)} \cdot \frac{f(w)}{F(w)} \cdot \frac{1}{F(w)^{N-1}}$. Therefore ${b_r^*}'(w)$ is decreasing (i.e. b_r^* is concave) in w and since ${b_r^*}'(w_r^*) = 1$, it follows that for all $w > w_r^*$, $b_r^*(w) < w$.

[17] See Foucault et al. (2013) for a recent survey.

2. Let $\hat{b}_w \in \arg\max_{b\in[0,w]} F(b)^{N-1}\left(\nu^*_{r-2} - b\right)$. We next establish a useful bound:

$$\begin{aligned}
\frac{U_{r-1}(w) - U_r(w)}{F(w)^{N-1}} &= \frac{\max_{z\in[0,w]} F(z)^{N-1}(\nu^*_{r-2} - z) - \max_{z\in[0,w]} F(z)^{N-1}(\nu^*_{r-1} - z)}{F(w)^{N-1}} \\
&\leq \frac{F(\hat{b}_w)^{N-1}(\nu^*_{r-2} - \hat{b}_w) - F(\hat{b}_w)^{N-1}(\nu^*_{r-1} - \hat{b}_w)}{F(w)^{N-1}} \\
&= \frac{F(\hat{b}_w)^{N-1}}{F(w)^{N-1}}\left[\nu^*_{r-2} - \nu^*_{r-1}\right] \\
&\leq \nu^*_{r-2} - \nu^*_{r-1}.
\end{aligned}$$

Using the preceding inequality,

$$\begin{aligned}
\nu^*_{r-2} - \nu^*_{r-1} &\geq \frac{U_{r-1}(w) - U_r(w)}{F(w)^{N-1}} \\
&= \frac{F(w)^{N-1}\left(\nu^*_{r-2} - b^*_{r-1}(w)\right) - F(w)^{N-1}\left(\nu^*_{r-1} - b^*_r(w)\right)}{F(w)^{N-1}}.
\end{aligned}$$

Rearranging terms, $\nu^*_{r-2} - \nu^*_{r-1} - b^*_{r-1}(w) + b^*_r(w) \leq \nu^*_{r-2} - \nu^*_{r-1} \implies b^*_r(w) \leq b^*_{r-1}(w)$.

□

Proof of Corollary 2.

1. By definition, $b^*_r(w) = \nu^*_{r-1} - \frac{U_r(w)}{F(w)^{N-1}} \leq \nu^*_{r-1}$. Since ν_{r-1} is a constant value, taking expectations of both sides gives $\int_0^{\bar{w}} b^*_r(w)g(w)dw \leq \nu^*_{r-1} \implies \nu^*_r \leq \nu^*_{r-1}$.

2. First consider the difference

$$b^*_r(w) - b^*_{r+1}(w) = \nu^*_{r-1} - \nu^*_r - \left(\frac{U_r(w) - U_{r+1}(w)}{F(w)^{N-1}}\right).$$

To conclude that $b^*_r(w) - b^*_{r+1}(w) \leq \nu^*_{r-1} - \nu^*_r$ it is sufficient to verify that $U_r(w) \geq U_{r+1}(w)$. To see this, recall that $b^*_{r-1}(w) \geq b^*_r(w)$. Hence, $\nu^*_{r-1} \geq \nu^*_r$. And so, $U_r(w) = \max_{b\in[0,w]} F(b)^{N-1}(\nu^*_{r-1} - b) \geq \max_{b\in[0,w]} F(b)^{N-1}(\nu^*_r - b) = U_{r+1}(w)$. Therefore, $b^*_r(w) - b^*_{r+1}(w) \leq \nu^*_{r-1} - \nu^*_r$. Taking expectations gives $\nu^*_r - \nu^*_{r+1} \leq \nu^*_{r-1} - \nu^*_r$ as required.

□

Proof of Corollary 3.

1. Follows directly from the proof of Corollary 1.

2. The expected payoff of a bidder of type w when following the strategy outline in Theorem 1 is

$$U_r(w) = \max_{b \in [0,w]} F(b)^{N-1}(\nu_{r-1} - b)$$

This expression is increasing in ν_{r-1}. Since $\nu_r < \nu_{r-1}$, $U_{r+1}(w) \leq U_r(w)$.

$\square$

References

Ausubel, L. M., Burkett, J. E., and Filiz-Ozbay, E. (2013). An experiment on auctions with endogenous budget constraints. Mimeo, University of Maryland.

Blume, L. E., Easley, D. A., Kleinberg, J., and Tardos, E. (2009). Trading networks with price-setting agents. *Games and Economic Behavior*, 67(1):36–50.

Bobkova, N. (2016). First price auction with asymmetrically budget constrained bidders. Mimeo.

Charness, G., Corominas-Bosch, M., and Fréchette, G. R. (2007). Bargaining and network strucutre: An experiment. *Jornal of Economic Theory*, 136:28–65.

Che, Y.-K. and Gale, I. (1996). Expected revenue of all-pay auctions and first price sealed-bid auctions with budget constraints. *Economics Letters*, 50(3):373–379.

Che, Y.-K. and Gale, I. (1998). Standard auctions with financially constrained bidders. *Review of Economic Studies*, 65(1):1–21.

Choi, S., Galeotti, A., and Goyal, S. (2014). Trading in networks: Theory and experiments. Working Paper 2014/08, Cambridge-INET Institute.

Condorelli, D. and Galeotti, A. (2014). Strategic models of intermediation networks. Mimeo.

Condorelli, D., Galeotti, A., and Renou, L. (2016). Bilateral trading in networks. *Review of Economic Studies*.

Foucault, T., Pagano, M., and Röell, A. (2013). *Market Liquidity: Theory, Evidence, and Policy*. Oxford University Press, New York.

Fox, J. (2002). Nonparametric regression. Appendix to *An R and S-PLUS Companion to Applied Regression*. Accessed February 20, 2015. `http://cran.r-project.org/doc/contrib/Fox-Companion/appendix-nonparametric-regression.pdf>`.

Gale, D. and Kariv, S. (2007). Financial networks. *American Economic Review: Papers and Proceedings*, 97(2):99–103.

Gale, D. M. and Kariv, S. (2009). Trading in networks: A normal form game experiment. *American Economic Journal: Microeconomics*, 1(2):114–132.

Jackson, M. O. (2008). *Social and Economic Networks*. Princeton University Press, Princeton, NJ.

Judd, J. S. and Kearns, M. (2008). Behavioral experiments in networked trade. In *EC '08: Proceedings of the 9th ACM Conference on Electronic Commerce*.

Kagel, J. H. (1995). Auctions: A survey of experimental research. In Kagel, J. H. and Roth, A. E., editors, *Handbook of Experimental Economics*. Princeton University Press, Princeton, NJ.

Kagel, J. H. and Levin, D. (2014). Auctions: A survey of experimental research. Mimeo.

Kosfeld, M. (2004). Economic networks in the laboratory. *Review of Network Economics*, 3(1):20–41.

Kotowski, M. H. (2010). First-price auctions with budget constraints: An experiment. Mimeo.

Kotowski, M. H. (2016). First-price auctions with budget constraints. Mimeo.

Kotowski, M. H. and Leister, C. M. (2014). Trading networks and equilibrium intermediation. Mimeo.

Kranton, R. E. and Minehart, D. F. (2001). A theory of buyer-seller networks. *American Economic Review*, 91(3):485–508.

Li, D. and Schürhoff, N. (2014). Dealer networks. Mimeo.

Manea, M. (2014). Intermediation in networks. Mimeo.

Milgrom, P. R. and Weber, R. J. (1985). Distributional strategies for games with incomplete information. *Mathematics of Operations Research*, 10(4):619–632.

Pitchik, C. and Schotter, A. (1988). Perfect equilibria in budget-constrained sequential auctions: An experimental study. *RAND Journal of Economics*, 19(3):363–388.

Rubinstein, A. and Wolinsky, A. (1987). Middlemen. *Quarterly Journal of Economics*, 102(3):581–594.

Smith, V. L. (1962). An experimental study of competitive market behavior. *Journal of Political Economy*, 70(2):111–137.

Smith, V. L. (1965). Experimental auction markets and the Walrasian hypothesis. *Journal of Political Economy*, 73(4):387–393.