Liquidity risk in sequential trading networks

Citation

Kariv, Shachar, Maciej H. Kotowski, and C. Matthew Leister. 2018. "Liquidity Risk in Sequential Trading Networks." *Games and Economic Behavior* 109 (May): 565–581. doi:10.1016/j.geb.2018.02.004.

Published Version

doi:10.1016/j.geb.2018.02.004

Permanent link

<http://nrs.harvard.edu/urn-3:HUL.InstRepos:35165081>

Terms of Use

This article was downloaded from Harvard University's DASH repository, and is made available under the terms and conditions applicable to Open Access Policy Articles, as set forth at <http://nrs.harvard.edu/urn-3:HUL.InstRepos:dash.current.terms-of-use#OAP>

Share Your Story

The Harvard community has made this article openly available.
Please share how this access benefits you. [Submit a story](#).

[Accessibility](#)

Liquidity Risk in Sequential Trading Networks*

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

October 9, 2017

Abstract

This paper studies a model of intermediated exchange with liquidity-constrained traders. Intermediaries are embedded in a trading network and their financial capacities are private information. We characterize our model’s monotone, pure-strategy equilibrium. Agents earn positive intermediation rents in equilibrium. An experimental investigation supports the model’s baseline predictions concerning agents’ strategies, price dynamics, and the division of surplus. While private financial constraints inject uncertainty into the trading environment, our experiment suggests they are also a behavioral speed-bump, preventing traders from experiencing excessive losses due to overbidding.

Keywords: Experiment, economic networks, intermediation, liquidity, auctions, budget constraints

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

In many decentralized markets, intermediary traders link producers and consumers. Two key variables determine the effectiveness of market intermediaries. First, a trader needs to have adequate financial resources—either in the form of cash on hand or credit—to pay for the goods that he wishes to trade. And second, each intermediary depends on his network of potential counter-parties to source goods and to find buyers. The interaction between these variables 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 and resold in the market?
3. What are the distributional consequences of the economy’s network structure? Are some traders systematically advantaged because of their relationships or position in the market?

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. We pursue complementary tracks involving theoretical and experimental analysis; we develop a tractable model that we test in the laboratory. Potential trading relationships and financial constraints are often shrouded or deliberately concealed in real-world markets. In a lab setting, we can directly control these key variables influencing behavior.

Our model, which builds on Gale and Kariv (2009), is simple, yet 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”).¹ 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. An instance of such a situation is illustrated in Figure 1, with four intermediary traders. A network specifies the relationships among the agents and determines the feasible transactions among them. If agent i is linked to agent 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 passage along links in the network from the seller to the buyer.

The key feature of our model, and where we substantively depart from Gale and Kariv (2009), is that each trader is liquidity or budget constrained. He has limited funds to finance

¹Following 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.

Figure 1: A trading network.

his trading activity. These constraints are private information and inject considerable risk into the market. A successful trader must anticipate the funds available to his immediate counter-parties and be mindful of similar constraints elsewhere in the economy. For simplicity, we assume that transaction prices are set by a first-price auction. In equilibrium, intermediaries in our economy adopt subtle bidding strategies accounting for the compounding financing risks. Average prices systematically rise as the asset nears the final buyer. Traders closer to the buyer benefit from the reduction of uncertainty. They manage to garner a relatively larger share of the surplus than traders further away.

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. Agents with ample funds shade their bid relative to their trading budget and intermediaries closer to the final buyer adopt uniformly more aggressive bidding strategies. Consequently, prices rise and intermediaries closer to the buyer tend to earn higher trading profits. Interestingly, our experiment also identifies a practical disciplinary role played by financial frictions and budget constraints. In the laboratory, we observe that traders who are flush with resources tend to experience a mild decline in profits relative to the equilibrium benchmark and to others. 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. The latter effect often dominates.

This paper is organized as follows. Section 2 briefly reviews the related literature. We link our study to both theoretical and experimental studies of networked markets and auctions. Section 3 introduces our model. The model’s equilibrium analysis is presented in Section 4. We derive several implications concerning intermediary behavior, price dynamics, and welfare. Our experiment’s procedures are summarized in Section 5. Data analysis is performed in Section 6. Section 7 provides a discussion of our results and our study’s broader implications. Section 8 concludes. All proofs are gathered in Section 9. An Online Appendix contains supplementary material pertaining to our study and experiment.

2 Related Literature

This paper contributes to literatures on networked markets and auctions. 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. (2017), Manea (forthcoming), and Condorelli et al. (2017). Condorelli and Galeotti (2015) 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.

As noted above, our analysis is closest to that of Gale and Kariv (2009), who study a similar class of trading economies. Assuming a complete information setting, Gale and Kariv (2009) show that prices converge to the asset’s value and all intermediation rents are quickly competed away. Our study differs in two important ways. Most importantly, we study an economy where intermediaries hold private information concerning their trading ability, which we model as a budget to bankroll transactions. Second, our price-setting mechanism, a standard first-price auction, is simpler than Gale and Kariv’s protocol, which involves averaging bid and ask prices.² These departures lead to predictions distinct from those advanced by Gale and Kariv (2009). Average prices in our economy rise with successive transactions and traders enjoy positive intermediation rents.

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. Choi et al. (2017) examine intermediated exchange theoretically and in the laboratory focusing on a posted-price mechanism. Their experiment highlights the importance of coordination among market intermediaries and the role of “critical traders.” Critical traders enjoy monopoly-like positions in a network allowing them to extract considerable intermediation rents. Our economy does not feature critical traders; instead, we identify an alternative channel allowing intermediaries to earn sustained profits in equilibrium. Kosfeld (2004) provides a survey of other laboratory experiments involving networked economies.

²Gale and Kariv (2009) report results from a control treatment where transaction prices equaled only the bid, as in our model. These results are similar to their baseline experiment with the more complex pricing rule. Mindful of this result, we have opted for a simpler pricing rule from the onset. See also Remark 2.

Though our focus is on trade in a networked economy, we also contribute to the experimental auction literature. Our experiments incorporate a direct test of Che and Gale’s (1996) model of a first-price auction with private budget constraints. In this model, each bidder has a one-dimensional type, a private budget constraint above which he cannot bid, and the item for purchase has a commonly-known and fixed value.³ To our knowledge, Che and Gale’s (1996) model has not been tested in a laboratory before. Anticipating our results, data from our baseline treatment are consistent with the equilibrium that they characterize. Other experimental investigations of auctions with budget constraints include Kotowski (2010) and Ausubel et al. (2017), 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 (2016).

3 The Model

In this section we introduce the model that will inform our experimental analysis. We characterize its equilibrium in the next section.

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. We focus on a tractable class of networks originally analyzed by Gale and Kariv (2009). Choi et al. (2017) call such networks “multipartite networks.” Figure 2 illustrates a family of multipartite trading networks involving increasing degrees of intermediation between the buyer and the seller. In a multipartite network, agents are arranged in rows $0, 1, \dots, R + 1$. 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). 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

³Most models of auctions with private budget constraints, such as Che and Gale (1998) or Kotowski (2017), assume bidders’ types are two-dimensional, encompassing private information concerning the item’s value and a private budget constraint. Bobkova (2017) generalizes Che and Gale’s (1996) model to an asymmetric setting.

Figure 2: A family of multipartite trading networks.

rows $1, \dots, R$. Rows are numbered according to their network distance from the buyer.⁴ Traders, who are risk neutral, do not value the asset per se. Rather they seek to earn profits by facilitating trades in the network. 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 $r - 1$, as illustrated in Figure 2 with the corresponding directed links. As usual, a trader earns profits if he can “buy low” and “(re)sell high.”

The particular trading networks illustrated in Figure 2 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. For simplicity, 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 2, 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 (by assumption) pays v for the good.

With the exception of the pricing rule, the model above broadly parallels that of Gale

⁴Our numbering scheme simplifies our equilibrium characterization. Gale and Kariv (2009) follow the opposite numbering convention.

and Kariv (2009). Substantially 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 each intermediary trader faces a private budget constraint that represents his access to funds or liquidity 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 trading budgets 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 non-decreasing in w . Many common distributions, including the uniform, 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. When there is one row of intermediary traders ($R = 1$), our model reduces to Che and Gale's (1996) model of a common-value, first-price, sealed-bid auction with private budget constraints.

4 Equilibrium Analysis and Model Predictions

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:

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

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

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

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). An immediate consequence of the equilibrium's inductive structure is that the strategy defined in Theorem 1 is independent of the parameter R . That is, if b_r^* is the equilibrium bidding strategy of a row- r trader in a network with $R \geq r$ rows (and N traders per row), then b_r^* continues to be the equilibrium bidding strategy of a row- r trader in a network with $R' \geq r$ rows (and N traders per row).

Before discussing the above equilibrium in greater detail, we present 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 between 0 and 100. Assume that $N = R = 3$, as in Figure 2(c). The equilibrium bidding strategies are illustrated in Figure 3. These bidding strategies are:⁵

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

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

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

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.

This example nests equilibria from other trading networks as special cases. In particular,

⁵When stating the strategy, we round several values for expositional ease.

Figure 3: Equilibrium bidding strategies in Example 1.

if $R = 1$ and $N = 3$, then the equilibrium strategy is composed only of b_1^* . If $R = 2$ and $N = 3$, then b_1^* and b_2^* characterize the traders' equilibrium bidding strategies.

Example 1 suggests several general characteristics 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. The first corollary focuses on two striking features of individual-level bidding strategies.

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 w , a trader positioned further from the buyer bids less, i.e. $b_r^*(w) \geq b_{r+1}^*(w)$.*

The conclusions of Corollary 1 are readily apparent in Example 1 and Figure 3. First, bidders' strategies are increasing and concave. For bidders in each row, there is a critical budget 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 decreasing fraction of his funds to the transaction.

Second, bidders' strategies are ordered uniformly. 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 of 100. A trader in row 2, however, can only resell the asset to a trader in row 1. Since the bids of row 1 traders are capped by their budget, 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.

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 budget. Thus, and following directly from Theorem 1, realized transaction prices 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, i.e. $\nu_{r+1}^* \leq \nu_r^*$.*
2. *The difference in expected prices is increasing as the asset approaches the buyer, i.e. $\nu_r^* - \nu_{r+1}^* \leq \nu_{r-1}^* - \nu_r^*$.*

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, price differences evolve in the opposite manner.

Among all intermediaries, traders closest to the buyer capture the greatest share of the expected total surplus. Of course, part of the total surplus goes to the seller as well.

As expected 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 row- r trader with a budget 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 a budget less than w and bid less than $b_r^*(w)$ in equilibrium. 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. Conditional on w , expected payoffs are increasing as the distance to the buyer decreases, i.e. $U_{r+1}(w) \leq U_r(w)$.

Corollary 3 shows that traders' type-contingent expected payoffs differ systematically through the network. Greater expected payoffs 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 contrast with the conclusions of Gale and Kariv (2009), who argue that intermediation rents are entirely competed away. In our setting, the presence of private information ensures that intermediation rents remain positive.

Remark 1. Our analysis and experiment focus on the model's symmetric, monotone, pure-strategy equilibrium, which is both intuitive and focal. Our model also admits many mixed-strategy and non-monotone equilibria. These equilibria replicate the bid distribution induced by the strategies defined in Theorem 1. Intuitively, since traders with high budgets are indifferent among a range of bids (Corollary 3), their bidding strategies can be "rearranged"

without upsetting the equilibrium provided the overall bid distribution remains unchanged. We elaborate on this point in Online Appendix A. Importantly, many of our theoretical conclusions depend only on the bid distribution; hence, they are robust across equilibria.

Remark 2. A further extension of our setup introduces reserve prices into the pricing protocol. Under this scheme, an agent currently selling the asset could specify a minimum acceptable bid. It is straightforward to extend the analysis of Che and Gale (1996), and thus Theorem 1, to allow for reserve prices. Generally, the optimal (and equilibrium) reserve price depends on the economy’s parameters. In Example 1, which parallels our experimental study, the revenue-maximizing reserve price is zero for all traders. The supporting intuition is as follows. Traders with low budgets “bid everything” in equilibrium and a small reserve price only screens them out from the auction. A large reserve price can potentially increase the sale price, but the reduction in sale probability dominates. Hence, expected payoffs decline. Given the optimality of a zero reserve price in our experiment, we have suppressed their role in our analysis and model.

4.1 A Practical Extension: Limited Liability

Several practical considerations complicate the direct translation of the preceding model to a laboratory setting. One difficulty 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 subsequently receives from a downstream counter-party. Standard experimental protocols and laboratory rules (including ours) do not allow a subject to incur a financial penalty in a study. Numerous methods have been employed in the literature to address this common concern, sometimes with mixed effects.⁶ A simple method, which we follow in our experiment, involves introducing a limited-liability constraint. Specifically, a trader can be protected from loss by computing his earnings as

$$\text{Earnings} = \max \left\{ \underbrace{\text{Resale Price} - \text{Purchase Price}}_{\text{Trading Profits}}, 0 \right\}. \quad (4)$$

Gale and Kariv (2009) compute earnings in their experiment with the same formula. They argue that (4) can be interpreted as a compensation scheme for a professional trader who enjoys bonuses when trading profits are positive, but earns constant wages otherwise.

⁶This concern also arises in studies of common-value auctions due to the winner’s curse. See Kagel (1995) for a survey.

When earnings are computed according to (4), traders' incentives change relative to our original model. Since agents are protected from a loss, intuition suggests that equilibrium bidding ought to be more aggressive. In Online Appendix B we provide an equilibrium characterization in our economy, analogous to Theorem 1, when (4) defines traders' payoffs. Though considerably less tractable than the baseline model, the equilibrium can again be defined inductively. Equilibrium bidding strategies can no longer be expressed in closed form; however, they can be readily computed numerically in standard examples.

To gauge the impact of limited liability, in Figure 4 we illustrate the equilibrium strategies from Example 1 when earnings are instead computed according to formula (4). Recall that the example featured three rows of traders, with three traders per row. Budgets were uniformly distributed between 0 and 100. The asset's value to the final buyer was 100. (These parameters are identical to those employed in our experiment.) In the figure, we illustrate the equilibrium strategies in the presence of limited liability as dashed curves and label them as $\tilde{b}_r(\cdot)$. For comparison, the figure also illustrates the equilibrium strategies from our baseline model without limited liability, denoted by $b_r^*(\cdot)$. The 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. 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. Paralleling our baseline analysis, for each row r we define $\tilde{w}_r$ as the critical budget level such that $\tilde{b}_r(w) = w$ for all $w < \tilde{w}_r$ and $\tilde{b}_r(w) < w$ for all $w > \tilde{w}_r$. With limited liability, the critical values ($\tilde{w}_r$) are 43.5, 50.5, and 67. Under its original specification, the critical values (w_r^*) in Example 1 are 39, 46, and 67.

Since limited liability affects bidding in a measured manner, the qualitative conclusions outlined in Corollaries 1–3 continue to obtain. The continued applicability of Corollary 1 is apparent in Figure 4. Given the example's parameterization, Corollaries 2–3 can be easily verified, either analytically or numerically. Since many of our original conclusions continue to apply, the limited liability constraint provides a practical, minimally invasive compromise in implementing our experimental market.

4.2 Predictions

We conclude our theoretical analysis by translating our formal results into a set of empirical predictions that will form the basis of our experimental investigation. These predictions draw on Theorem 1, Corollaries 1–3, and our extension to the case of limited liability.

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.

First, monotone equilibria are focal in many auction-like games; thus, we anticipate traders' bidding strategies to exhibit the following qualitative characteristics.

Prediction 1 (Monotone Bidding).

1. Within a row, traders with greater budgets will place greater bids; however, traders with large budgets will not expend all of their funds.
2. Conditional on budget level, traders closer to the buyer will exhibit more aggressive bidding behavior.

Second, Corollary 2 describes the evolution of prices in the economy. Particular prices are independent of each other, but higher average prices should be observed in rows closer to the buyer.

Table 1: Experimental sessions and treatments.

Network Type	Number of Sessions	Total Subjects	Decision Observations	Average Earnings (\$) *
1×3	1	30	1,500	26.04
2×3	2	72	3,600	24.51
3×3	3	90	4,500	24.03

* Reported earnings include the \$5 show-up fee and the \$10 participation fee paid at the experiment’s conclusion.

Prediction 2 (Prices). Average 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 when trading budgets are low and constant when they are high.
2. Conditional on his budget, a trader closer to the buyer earns 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 noted in Remark 1 above. Finally, as explained above, these predictions continue to apply even if earnings are subject to a limited liability constraint.

5 Experimental Design and Procedures

In our experiment, we consider a family of networks to test our three 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 budgets without introducing the confound of network depth. 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. Each subject received a payment of \$5 for showing up at the experiment on time. At the experiment’s conclusion, each subject received an additional payment of \$10 for participating in and completing the experiment plus their within-experiment earnings or trading profits. We describe the calculation of trading profits below. Total compensation for participation in the experiment ranged from \$15 (the subject earned only the show-up and participation fees) to \$87.90 (the subject earned the show-up and participation fees plus \$72.90 in trading profit). Within the experiment, we refer to units of currency using the neutral term “tokens.”

Each of six sessions consisted of 50 independent trading rounds. 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 budget for that round. This budget, called the “token endowment” within the experimental platform, was randomly and uniformly

Figure 5: Sample interface in the 3×3 network treatment.

drawn between 0 and 100 tokens (including decimal values) and not carried over between rounds. Budgets 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 budget 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 5 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 is 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 budget was displayed in the Token Endowment field, immediately above the Bid field. After each round of bidding, each subject was informed of others’ bids and budgets. They were also shown the sequence of executed trades and the associated prices, as illustrated in Figure 5. 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. Each round had an equal probability of being chosen. The subjects were then paid based on that round’s earnings. To compute final payments, tokens were converted into dollars at a 1-for-1 exchange rate. If a subject did not trade in the selected round, his or her trading profits were zero. The subject received only the \$5 show-up fee and \$10 participation fee. If a subject did trade in the selected round, he or she received the show-up and participation fees plus any earnings from the selected round. These earnings were computed according to formula (4), which incorporated the limited liability constraint. As explained above, use of this payment scheme ensured subjects did not incur a financial penalty; it was also used by Gale and Kariv (2009).⁷ Subjects were informed of the earnings formula as part of the experiment’s instructions.

6 Data Analysis and Results

As a starting point for our data analysis, we examine bidder behavior. In Figure 6 we plot LOESS curves⁸ (and 95% confidence bounds) summarizing average bids submitted by

⁷An 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. There is a risk of confusion between the buffer budget and the particular budget constraint in a particular round. Furthermore, loss-aversion may become more salient as a behavioral confound. We believe the present design is more closely aligned with our original model and we can account for the payoff formula’s equilibrium consequences directly, as explained in Section 4.1.

⁸LOESS is a standard non-parametric regression procedure, which fits a local polynomial regression to (in our case) bivariate data. See Fox (2008) for a practical discussion of this and related methods.

traders in each row as a function of their budget.⁹ We group our estimates by network type, but we pool across sessions and bidders.¹⁰ 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 it is low. High budget traders do not commit all of their funds to the auction. Instead, they shade their bid 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 estimates are necessarily less precise, due to fewer observations.

Result 1 (Bidding). Across networks, traders’ bids are increasing in their budget level, even when budgets are not a binding constraint. Traders with low budgets bid all of their available funds to acquire the asset. Traders with ample funds reduce their bid relative to 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

⁹All of our estimates and data analysis use the full data sample. As a robustness check, we have repeated our estimates dropping the first 10 and 25 rounds from each session to account for the possibility of learning. The results of this exercise were not substantively different from the estimates based on the full sample. Our results support Gale and Kariv’s (2009) conclusion that subjects appear to adapt quickly to this type of trading environment.

¹⁰In 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 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 6: Estimated bidding strategies. LOESS curves with 95% confidence bounds. All plots were created with the *R* statistical package using default settings (the span parameter is set to the default 0.75).

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

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.

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 of 100. 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 2.4 between rows 2 and 3.¹¹ 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 a trading budget, and across budget levels, conditional on row. To illustrate both comparative statics, in Figure 7 we plot traders’ payoffs as a function of their realized budget 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 budget level. Payoffs are computed accounting for the limited liability constraint as given by formula (4) and explained in Section 5. 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.

¹¹The median difference in realized prices differs from the difference in median realized prices, which can

(a) All networks.

(b) 1×3 Network.

(c) 2×3 Network.

(d) 3×3 Network.

Figure 7: Average earnings as a function of a trader's budget. LOESS curves with 95% confidence bounds. All plots were created with the *R* statistical package using default settings (the span parameter is set to the default 0.75).

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 at low budget levels. Traders with low budgets always have a small chance of successfully acquiring the asset; hence, payoffs are (unsurprisingly) uniformly low. When budgets are high, 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 in all panels of Figure 7. 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) and remains a persistent feature at finer scales.

To evaluate the significance of 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 budget 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 7, panel (a), and Table 4 corresponds to Figure 7, 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, an increase in a trader's budget is associated with a decline in average earnings. For example, increasing the budget of a row 1 trader from 67 to 100, thereby removing all constraints on his bid, reduces his average payoff by about 3 tokens (dollars). 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.¹² 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.

be computed directly from Table 2.

¹²Neglecting the limited liability constraint maintains the payoff ordering but significantly reinforces the downward trend in the payoffs of row 2 and row 3 traders with high budgets. Therefore, the correction for earnings formula (4) is important.

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

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

Robust standard errors in parenthesis. * $p < 0.1$. ** $p < 0.05$. *** $p < 0.01$.

Table 4: Earnings trends among high-budget traders by row and network. OLS regression with subject fixed effects. Pooled sample across sessions.

Network	1×3	2×3		3×3		
Row	1	1	2	1	2	3
Constant	24.996*** (6.149)	25.301*** (5.570)	7.932*** (3.244)	17.077*** (5.794)	10.244*** (2.901)	10.834*** (2.589)
Budget	-0.120** (0.058)	-0.137** (0.059)	-0.030 (0.027)	-0.031 (0.062)	-0.055** (0.025)	-0.039** (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. * $p < 0.1$. ** $p < 0.05$. *** $p < 0.01$.

We can rationalize the observed downward trend in payoffs with a simple behavioral argument. In our model, the blessing of a high budget comes with the curse of more opportunities for error. An intermediary with a tight budget constraint faces a relatively straightforward problem in formulating his bid. With high probability the asset’s resale value will exceed his budget. 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 trader is more complex. Agents 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 initially increasing in a trader’s budget. When a trader’s budget is relatively large, a further increase in his budget is associated with a mild, though statistically significant, decline in earnings.

7 Discussion and Interpretation

As with all market experiments, we build our analysis around a stripped-down theoretical model that we subsequently test with a particular population. Neither the trading platform nor the university-student subject pool perfectly replicates the “real world.” Of course, concerns regarding an experiment’s external validity are not new; see Chapters 14–20 of Fréchette and Schotter (2015) for an extensive survey. Our analysis studies general principles of market operation. Thus, echoing Kessler and Vesterlund (2015), we believe the qualitative insights of our study will generalize to related non-laboratory environments even though the quantitative estimates are specific to our model and experiment. Generalizable insights from our analysis include (i) the coordination on an equilibrium that is monotone in financial capacity,¹³ (ii) the payoff advantage enjoyed by intermediaries closer to a final buyer or

¹³We note, additionally, that equilibria are monotone *within* each transaction episode (i.e. auction) and *across* transactions occurring within the same market.

consumer, and (iii) the disciplinary benefit of budget constraints observed in the laboratory, which we interpret as preventing behavioral errors.¹⁴

Below we explain how two elements of our model, the multipartite network structure and the private budget constraints, relate to features beyond the laboratory environment.

Multipartite Trading Networks

A characteristic of our setting is the network structure. 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.

Beyond OTC markets, our network structure parallels some of the directed intermediation observed in financial systems and supply chains (Spulber, 1996). As large institutions save and consumers borrow, for example, financial intermediaries, i.e. banks and dealers, facilitate the movement of funds between wholesale and retail markets. Similarly, buyer financing is a prominent concern in many export industries, which often depend on intermediaries to fund complex transactions and to facilitate the movement of physical goods. Though stylized, we believe our model provides some insight to the strategic dynamics at play in these applications.

Multipartite networks also offer insights regarding more complex economies by focusing on local interactions among agents. They can be interpreted as a small neighborhood within a larger economy where a few agents enjoy a dense set of relationships among each other.

¹⁴This behavioral disciplinary effect complements the use of budget constraints to solve a principal-agent problem in an auction environment (Burkett, 2016; Ausubel et al., 2017).

Under this paradigm, the non-strategic seller and the buyer can be interpreted as metaphors for (not modeled) 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 also reflect an economic opportunity-cost of funds.

More abstractly, a reinterpreted private budget may capture behavioral constraints not directly related to financial capacity. For example, it may capture a trader’s relative status in the network. A trader with perceptions of weak ties to resale markets—represented by a low realized “budget”—may be unwilling or unable to commit a large sum to a particular interaction. In our economy, such low-budget traders tend to earn relatively large profits conditional on trading successfully; however, they trade with relatively low frequency.

8 Conclusion

We study a decentralized market with financing risk and a multipartite network structure, focusing on intermediary behavior. In our economy, (average) transaction prices rise with successive transactions and intermediaries positioned closer to the buyer enjoy greater expected profits. We find strong evidence supporting these theoretical predictions in our laboratory experiments. In the lab, we also identify a mild negative correlation between expected profits and subjects’ trading budgets, conditional on budgets being relatively high. That is, liquidity-rich traders tend to overbid. This trend is observed across rows. Tighter budget constraints appear to mitigate this behavioral tendency. Thus, budgets have a disciplinary role in markets, preventing excessively costly “trembles” or “errors.”

While the implication of private liquidity risk has been well-studied in centralized exchanges,¹⁵ 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. Analyzing

¹⁵See Foucault et al. (2013) for a recent survey.

specific policy interventions, such as liquidity injections that target trading positions within the market, offers a promising avenue for future research.

9 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. $\square$

Proof of Corollary 1.

1. Since $F(w)/f(w)$ is increasing, $F(w)^{N-1}(\nu_{r-1}^* - w)$ 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 $w > w_r^*$, $b_r^*(w) < w$.

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

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

Using the preceding inequality,

$$\begin{aligned}\nu_{r-2}^* - \nu_{r-1}^* &\geq \frac{U_{r-1}(w) - U_r(w)}{F(w)^{N-1}} \\ &= \frac{F(w)^{N-1} (\nu_{r-2}^* - b_{r-1}^*(w)) - F(w)^{N-1} (\nu_{r-1}^* - b_r^*(w))}{F(w)^{N-1}}.\end{aligned}$$

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 $\nu_r < \nu_{r-1}$, $U_{r+1}(w) \leq U_r(w)$.

□

References

- Ausubel, L. M., Burkett, J. E., and Filiz-Ozbay, E. (2017). An experiment on auctions with endogenous budget constraints. *Experimental Economics*.
- 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. (2017). Asymmetric budget constraints in a first price auction. Mimeo.
- Burkett, J. (2016). Optimally constraining a bidder using a simple budget. *Theoretical Economics*, 11:133–155.
- Charness, G., Corominas-Bosch, M., and Fréchette, G. R. (2007). Bargaining and network structure: An experiment. *Journal 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. (2017). Trading in networks: Theory and experiments. *Journal of the European Economic Association*, 15(4):784–817.
- Condorelli, D. and Galeotti, A. (2015). Strategic models of intermediation networks. Mimeo.
- Condorelli, D., Galeotti, A., and Renou, L. (2017). Bilateral trading in networks. *Review of Economic Studies*, 84(1):82–105.
- Foucalt, T., Pagano, M., and Röell, A. (2013). *Market Liquidity: Theory, Evidence, and Policy*. Oxford University Press, New York.
- Fox, J. (2008). *Applied regression analysis and generalized linear models*. Sage, Los Angeles.
- Fréchette, G. R. and Schotter, A., editors (2015). *Handbook of Experimental Economic Methodology*. Oxford University Press.
- Gale, D. M. 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.
- 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. (2016). Auctions: A survey of experimental research. In Kagel, J. H. and Roth, A. E., editors, *The Handbook of Experimental Economics*, volume 2, chapter 9, pages 563–637. Princeton University Press, Princeton, NJ.
- Kessler, J. B. and Vesterlund, L. (2015). The external validity of laboratory experiments: The misleading emphasis on quantitative effects. In Fréchette, G. R. and Schotter, A., editors, *Handbook of Experimental Economic Methodology*, chapter 18. Oxford University Press.
- 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. (2017). 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. (Forthcoming). Intermediation and resale in networks. *Journal of Political Economy*.
- Pitchik, C. and Schotter, A. (1988). Perfect equilibria in budget-constrained sequential auctions: An experimental study. *RAND Journal of Economics*, 19(3):363–388.
- Spulber, D. F. (1996). Market microstructure and intermediation. *Journal of Economic Perspectives*, 10(3):135–152.