Trading Networks and Equilibrium Intermediation*

Maciej H. Kotowski[†] C. Matthew Leister[‡]

June 10, 2014

Abstract

We consider a network of intermediaries facilitating exchange between buyers and a seller. Intermediary traders face a private trading cost, a network characterizes the set of feasible transactions, and an auction mechanism sets prices. We examine stable and equilibrium networks. Stable networks, which are robust to agents' collusive actions, exist when cost uncertainty is acute and multiple, independent trading relationships are valuable. A free-entry process governs the formation of equilibrium networks. Such networks feature too few intermediaries relative to the optimal market organization and they exhibit an asymmetric structure amplifying the shocks experienced by key intermediaries. Welfare and empirical implications of stable and equilibrium networks are investigated.

Keywords: Networks, Intermediation, Trade, Network Formation, Second-Price Auction, Supply Chains, Financial Networks
JEL: D85, L14, D44

*We thank Shachar Kariv and numerous seminar and conference audiences for valuable feedback. Leister thanks Sanjeev Goyal and the Economics Faculty at Cambridge University for hospitality. We are grateful for the referees and the editor for valuable suggestions that challenged us to (successfully) generalize several preliminary results. First version: February 29, 2012.

[†]John F. Kennedy School of Government, Harvard University, 79 JFK Street, Cambridge MA 02138. E-mail: <maciej_kotowski@hks.harvard.edu>

[‡]Department of Economics, 530 Evans Hall, University of California, Berkeley CA 94720-3880. E-mail: <leister@econ.berkeley.edu>

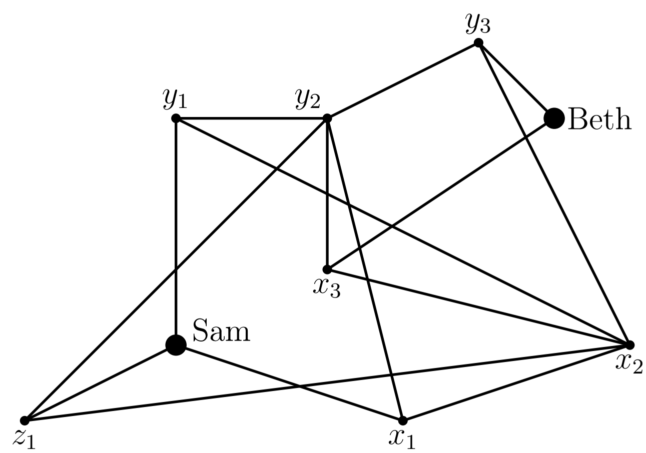

Figure 1: A trading network.

Intermediation in markets is commonplace. Consider the situation depicted in Figure 1. Sam is a farmer growing watermelons in California while Beth is a consumer of watermelons in New England. There are gains from trade between Sam and Beth; however, rarely will Sam and Beth trade directly. They likely do not even know each other. Instead, trade between them is mediated by a network of intermediary agents $\{x_1, x_2, x_3, y_1, y_2, y_3, z_1\}$. These intermediaries—such as wholesalers, transporters, distributors, or retailers—have invested in market-specific technologies and have developed a web of trading relationships through which Sam and Beth are linked. There are many paths in the economy through which Sam's produce can arrive on Beth's picnic table.

Intermediary networks like those in Figure 1 embody two cross-cutting features. First, there is *competition* among intermediaries with similar links and relationships. Intermediaries with overlapping relationships—like x_3 and y_3 in Figure 1—can perform similar tasks in the market and will compete to offer their services. Intuitively, overlapping relationships enhance a market's robustness. A shock experienced by a particular agent is unlikely to harm the market as a whole. Other agents, with similar relationships, can act as close substitutes ensuring goods continue to flow. Second, there is *complementarity* among intermediaries with dissimilar links and relationships. An intermediary who is close to a final consumer, like x_3, relies on intermediaries near a producer (x_1, y_1, or z_1) to kick-off the intermediation chain. Similarly, an intermediary who is close to a producer relies on those with direct links to final consumers to channel demand.

In light of the competitive and complementary forces embedded in a networked economy, two questions naturally arise.

1. What economic incentives sustain the arms-length nature of trading relationships in a network despite the presence of both complementarities and competition? Agents' desires to capitalize on complementarities and to constrain competition creates incentives for collusion or mergers. If this happens, the trading network is altered; therefore, its initial configuration lacked persistence and was unstable.

2. Are intermediary networks predisposed to adopt a form that reinforces market robustness or a form that begets market fragility? The push and pull of competition and complementarity suggests that either outcome is plausible a priori.

To answer these questions we propose a model of intermediated trade in a networked market. Our model is lean to maximize interpretive flexibility and it melds four classic ideas beyond the underlying network structure: (1) intermediaries have private trading costs, (2) an auction mechanism sets prices, (3) a core-like notion defines network persistence and stability, and (4) a free-entry/zero-profit condition drives network formation.

In our model, one good ("the asset") is traded and final consumers ("buyers") are separated from the good's producer ("the seller") by several tiers of intermediaries. In each tier, traders compete to provide intermediation services. Each intermediary bids to acquire the tradable asset with the aim of reselling it at a profit to neighbors, who in turn do the same until the asset is consumed by a final buyer. Though the network structure is common knowledge, each trader's private trading cost introduces residual uncertainty regarding intermediaries' demand into the environment. When a trader experiences a negative cost shock, the market's operation is shaken. However, if the web of relationships among intermediaries is sufficiently dense, such shocks minimally impact the market as a whole. If the trading network is locally sparse, a shock's impact is exaggerated and market breakdown may ensue.

Paralleling questions 1 and 2, we distinguish between stable and equilibrium trading networks. In our analysis, "stability" refers to a network's persistence and is distinct from the network formation process discussed above. Stable networks are immune to the contractive incentives implicit in networked markets and preserve traders' arms-length interactions. In a stable market, existing traders must not be able to profitably merge together in an attempt to exploit complementarities or to curtail competition. Our stability notion captures this intuition and allows us to isolate the distinct, and sometimes subtle, channels through which the incentives to destabilize an existing network operate. Our model suggests that the gain from curtailing competition (collusion among similar traders) is often greater than the gain from enhanced scope (collusion among dissimilar or complementary traders). While direct competitors unambiguously hurt a trader's profits, maintaining relationships with multiple independent complementary agents adds a valuable layer of robustness, which challenges the benefits otherwise associated with scope economies. Accordingly, we show that a network is stable when agents are subject to sufficiently frequent shocks, as then the benefits of multiple independent trading partners outweigh the net benefits of collusion. By stabilizing the network, idiosyncratic risk acts as a countervailing force to collusion. Thus, it helps preserve a relatively more efficient (competitive) market organization.

Equilibrium networks are the result of a network formation process, which we assume is governed by the free entry of intermediaries into distinct, specialized roles. Our model shows that this process results in networks featuring too few intermediaries relative to a socially-optimal network organization. Moreover, these few intermediaries additionally assume a configuration that exaggerates the negative shocks experienced by some traders. These conclusions spring from a fundamental wedge between the private incentives of an intermediary to operate in a market and the social benefit generated by that intermediary's activity. By adding a new path for the flow of goods, an intermediary competes with some traders but complements others. The unappropriated benefits from complementarity are sufficiently strong to result in an under-provision of intermediary services.

Though we offer reinterpretations of our model with an eye toward production and financial intermediation (see section 1), our model has a simple interpretation as describing a supply chain, like in the vignette above. In this case, equilibrium networks accord-well with many common empirical features of supply chains. For example, we show that in equilibrium there are more intermediaries near consumers ("retailers") than there are intermediaries near producers ("wholesalers").[1] This result is driven by asymmetries in the degree of complementarity among traders and in the uncertainty experienced by traders in different parts of the economy. It arises despite the absence of scale economies. Similarly, well-known empirical features of supply-chain networks, such as the "bull-whip effect" (Lee et al., 1997a,b), are easily discernible in the equilibrium networks of our model.

We develop our argument progressively. Section 1 introduces our model. Throughout, we take the presence of intermediaries as given and we focus on the interactions among

[1] In other words, our model suggests many retailers will carry similar items and these will be sourced from a smaller pool of wholesalers. Large multi-product retailers do not arise as there is only one good in our economy.

them.[2] In section 2 we examine price-formation and exchange taking the network structure as given. Section 3 considers network stability and we propose our model of network formation in section 4. Section 5 investigates the relationships between stable and equilibrium networks. We note parallels between our analysis and other studies as they arise and we link our conclusions to the wider literature on networked markets before concluding in section 6. In that section we also outline extensions and variations of our basic model. For example, throughout we assume that traders face uncertain *demand* from intermediaries for the asset. A "reversal" of our model accommodates *supply* uncertainty with parallel conclusions. Appendix A collects proofs. Appendices B and C are available in an online supplement.

1 Model

An economy is characterized by three elements. First, agents are organized in a network defining trading possibilities. Second, each trader has a private trading cost determining the prudence of exchange. Finally, a trading protocol sets prices. After introducing our model, we comment on our assumptions and we offer interpretations in relation to the exchange of goods, to production with intermediate inputs, and to financial intermediation.

Trading Possibilities Trading possibilities are summarized by a network. Agents are nodes while edges denote trading links. Our network topology generalizes the trading networks analyzed by Gale and Kariv (2009). Figure 2 presents a typical example. Agents are arranged in rows $0, 1, \ldots, R+1$, and trading possibilities conform to the following principle:

> An agent in row r can trade with any agent in rows $r+1$, r, and $r-1$ and vice-versa. Other trades for an agent in row r are not feasible.

This principle implies a lattice-like network as illustrated by Figure 2 for the case of $R = 2$.[3]

There are three types of agents in our economy. Row $R+1$ is inhabited only by the *seller*. The seller is the originator of a tradable asset that he is willing to sell at a price normalized to zero. Row 0 is inhabited by $n_0 \geq 2$ *buyers*. (Without loss of generality, we adopt the

[2] We do not model the underlying reason a particular market features intermediaries, though such arrangements are undoubtably common. Spulber (1996) argues that intermediation, broadly interpreted, accounts for a quarter of U.S. economic output. Intermediaries play important roles in markets with asymmetric information or search frictions. Moreover, legal constraints or natural barriers ensure that intermediaries become active market participants.

[3] Our numbering convention is opposite to the convention followed by Gale and Kariv (2009). Choi et al. (2014) examine a related class of networks, which they call "multipartite networks."

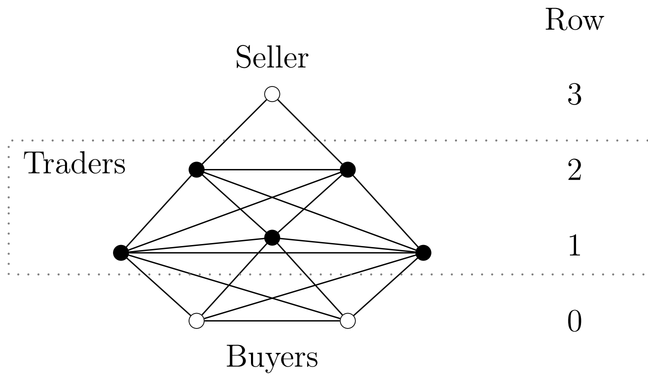

Figure 2: A trading network with configuration $\mathbf{n} = (3, 2)$.

convention that there are two buyers in all figures.) Each buyer is willing to pay $v > 0$ for the seller's asset. There exist gains from transferring the asset from the seller to a buyer; however, we assume that they cannot trade directly. Instead, there is a set of intermediary *traders* in rows $1, \ldots, R$ who may buy and (re)sell the asset. Traders do not value the asset per se; rather, they want to earn trading profits by executing network-conforming trades. Though we do not impose any ex ante restrictions on the direction(s) of trade, for ease of discussion we say row r is downstream (upstream) of row r' if $r < r'$ ($r > r'$).

Although our model features three types of agents, our main focus is on traders.[4] We let n_r be the number of traders in row r and we call $\mathbf{n} = (n_1, \ldots, n_R)$ the configuration of intermediary traders.[5] For example, in Figure 2 $\mathbf{n} = (3, 2)$. $\mathbf{n}$ captures two important characteristics of intermediary networks. Its length (R) measures the degree of intermediation in the economy. It might measure physical distance or it might summarize discrete steps in a supply chain. The number of traders in each row (n_r) measures the degree of competition among intermediaries. Agents in the same row are socially similar since their relationships with others overlap.

Trading Costs Each trader has a private trading cost $\theta_i \in \{0, t\}$, $0 < v < t$. A trader incurs the cost θ_i only upon acquiring the asset from another agent, even if he resells it. Intuitively, trading costs behave like an inventory cost but their interpretation can be broader.

[4]The "seller" and "buyers" can also be metaphors for larger (not-modeled) upstream and downstream markets.

[5]When needed, we use standard shorthand: $\mathbf{n}_{-r} = (n_1, \ldots, n_{r-1}, n_{r+1}, \ldots, n_R)$ and $\mathbf{n} = (n_r, \mathbf{n}_{-r})$.

For example, θ_i might capture marketing costs associated with resale. While trading costs are private, their distribution is common knowledge. Costs are distributed independently such that $\Pr[\theta_i = 0] = p$ for all i. We interpret p as describing the trading technology. If p is low, traders are often exposed to adverse cost shocks and trade is difficult. Although the economy's network structure is known, private trading costs imply that agents hold residual uncertainty concerning the liquidity of the (acquired) asset. Neighbors may or may not be willing to trade. By assumption, buyers and the seller have a trading cost of zero.[6]

Trading Protocol We assume that trade occurs via sequential second-price, sealed-bid auctions according to the following timeline.

0. Each agent learns his private trading costs.
1. When an agent holds the asset, he organizes an auction to sell it. Each of his neighbors in the network submits a bid from the set $\mathcal{B} = \{\ell\} \cup \mathbb{R}_+$.
   (a) The bid $\ell < 0$ is a non-competitive bid equivalent to not participating in the auction. If all auction participants bid ℓ, the asset is not sold and it expires. In this case, trade "breaks down."[7]
   (b) All bids $b \neq \ell$ are competitive bids. The agent submitting the highest competitive bid wins the auction. A uniform lottery resolves ties.
2. The agent winning the auction takes ownership of the asset and incurs his private trading cost. He makes a payment equal to the second-highest competitive bid (or zero if all others bid ℓ) to the auction's organizer.
3. Steps 1 and 2 repeat until trade breaks down or the asset reaches a buyer, who consumes it.

Traders are risk-neutral and wish to maximize trading profits (payments from others minus payments to others) net of trading costs. A buyer receives a payoff of v minus his payment.

[6] Allowing the seller and the buyers to incur trading costs does not substantively alter our results but complicates exposition.

[7] Though "breakdown" has an admittedly extreme connotation, we intend for the term to describe any event that interrupts the flow of goods in a manner adversely affecting their value. For example, it may correspond to a delay in the good's delivery. A buyer's valuation for a delayed product may be but a fraction of its original value. Similarly, in a production network it may correspond to the unavailability of an intermediate input good.

1.1 Discussion and Interpretation

Elements of our model deserve comment and elaboration. Crucial to our analysis is the incomplete, yet regularized, network structure. It aims to capture the two fundamental dimensions of intermediary markets. First, trading networks imply complementaries among intermediaries. Traders in different rows rely on each other to supply trading opportunities. If trade breaks down prematurely, downstream traders suffer. If expected downstream terms-of-trade deteriorate, upstream traders' expected profits fall as their expected resale values are impacted. Second, traders with similar relationships, i.e. those in the same row, are substitutes and competition drives their interaction. Our networks provide sufficient flexibility to examine both effects, which are also present in less-regularized markets.

To minimize confounds, we assume that all traders are ex ante identical, except for their position in the network. We can introduce some asymmetries—such as (small) row-dependent variation in p—without qualitatively changing our main conclusions. Similarly, if trading costs were continuously distributed the qualitative behavior of the model will be unchanged. Such a modification entails some quantitative amendments that render the exposition more involved.

Like Kranton and Minehart (2001) or Patil (2011), we rely on a second-price (equivalently, an ascending) auction to structure exchange. Beyond capturing the flavor of a competitive bidding process, this format allows us to bracket price-setting and to move quickly into a discussion of equilibrium and stable networks. Our analysis is robust to alternative pricing protocols provided revenue equivalence with the second-price auction obtains.[8] We leave to future research the explicit modeling of other trading or pricing schemes such as consignment (Rubinstein and Wolinsky, 1987), bargaining (Corominas-Bosch, 2004; Manea, 2011; Elliott, 2013; Condorelli and Galeotti, 2012a; Siedlarek, 2012), or posted prices (Choi et al., 2014).[9] Manea (2014) develops a model with a similar motivation to our environment. He relies on a bargaining game to set prices, assumes a directed graph to describe trading possibilities, and assumes that traders do not have a private trading cost.

In the introduction, we sketched our model's interpretation concerning the exchange of goods. In such an application, a network of intermediaries acts as a geographic and temporal bridge between producers and consumers. Another interpretation considers production with

[8] In an early draft (January 2012) we developed the model herein with first-price, sealed-bid auctions as the price-setting mechanism. Revenue equivalence obtains and our results continue to apply. In that model, traders place bids according to a mixed strategy in equilibrium, complicating exposition.

[9] Though we employ an auction mechanism to set prices, we do not optimize this mechanism. Therefore, incorporating an optimal auction (Myerson, 1981) with resale (Zheng, 2002) is a possible generalization.

intermediate inputs. A consumer wishes to purchase one unit of a final good at a price of v. Only firms in row 1 can produce this good. The good's production function combines one unit of labor (for example), at cost θ_i, with one unit of an intermediate good produced by firms in row 2; and so on. Interpreted in this way, our model emphasizes complementaries in production—a theme explored extensively in the literature on economic development (Kremer, 1993)—and competition among intermediate goods producers.[10]

Our model can also be viewed as a financial market. An investor (the seller) has one unit of capital available. A safe asset offers a return normalized to zero. Each firm seeking financing (a buyer) offers an expected net return of $v > 0$ for the funds. Intermediary financial institutions—banks, brokers, insurance companies, mutual funds, etc.—link the investor and the firms. The investor initially allocates his funds with a nearby, trusted intermediary promising the highest return. The intermediary does the same, and so on until the funds reach a firm. Intermediaries skim small fractions of the expected return promised by the firm as a payment for their intermediation services. Gofman (2011) proposes a model of an economic network to study financial transactions. Echoing elements of our model, he too assumes a single asset is traded and agents' valuations are private information. In contrast to our analysis, however, he employs a bargaining model to specify the trading mechanism and price formation.

As clear from the model, we take the need for intermediaries as a fundamental feature of the market under study. Thus, we do not explore the reasons for intermediation, which may include legal restrictions, technological specialization, or information imperfections. Rather, like Gale and Kariv (2009), Manea (2014), Wright and Wong (2014), or Choi et al. (2014)—to note but a few recent examples—we examine the operation of a market taking intermediation as given. By not specifying the underlying reason for intermediation, our model can be applied to any situation where successive intermediaries link buyers and sellers.

2 Exchange in a Fixed Network

We begin by studying trade in a fixed network. As our model embeds multiple second-price auctions, it necessarily admits multiple equilibria. Following tradition, we focus on an equilibrium where agents "bid their value" for the asset and it moves systematically towards buyers in row zero.

[10]Nagurney and Qiang (2009) employ networks with similar structures to ours to describe production inside a firm.

Theorem 1. *There exists a perfect Bayesian equilibrium of the trading game where each agent i (in row r) adopts the following strategy, denoted σ_i^*:*

1. *If the agent's costs are low and the asset is being sold by an agent in row $r+1$, the agent places a bid equal to the asset's expected resale value conditional on all available information and on σ_{-i}^*. (Buyers in row 0 bid their value, v, for the asset.)*

2. *Otherwise, the agent bids ℓ.*

Via inductive reasoning, using the buyers' bids as the anchor, the strategy profile proposed in Theorem 1 specifies a bid in all contingencies for every agent. Specifically, equilibrium-path expected resale values (bids) are defined inductively given the anticipated behavior of downstream traders. If

$$\delta(n) \equiv 1-(1-p)^n - np(1-p)^{n-1}, \tag{1}$$

then the asset's equilibrium-path expected resale value to a row-r trader is $\nu_r = \delta(n_{r-1})\nu_{r-1} = \prod_{k=1}^{r-1}\delta(n_k)v$. $\delta(n)$ is the probability that at least 2 out of n agents have a low trading cost. Only when there are 2 low-cost traders does the asset trade at a non-zero price.

Example 1. Suppose $\mathbf{n} = (3,2)$, as in Figure 2. Let $p = 1/2$ and $v = 1$. On the equilibrium path, low-cost row-1 traders bid 1. Low-cost row-2 traders bid $\delta(3)\cdot 1 = 1/2$.

We focus on the equilibrium in Theorem 1 due to its intuitive appeal and its reassuring characteristics. First, the asset does not "backtrack" nor does it pause and restart.[11] Trade has a natural direction toward the buyers. Second, expected prices and bids are nondecreasing as the asset approaches the buyers. Finally, (1) is increasing in p and n. Hence, average prices increase as low-cost traders become more common ($p\uparrow$) and as trader competition intensifies ($n_r\uparrow$).

Although intuitive, the "bid your expected resale value" strategy demands a high degree of trader sophistication. It is not a dominant strategy as it depends on others' anticipated behavior feeding into expected resale values. Traders must anticipate others' equilibrium bids and an error in the requisite inductive reasoning can compromise the outcome. Laboratory experiments studying a similar trading environments by Gale and Kariv (2009) and Gale et al. (2012) suggest that equilibrium predictions are in accord with observed outcomes. Whereas those studies do not investigate our exact trading game, we view their results as supporting our analysis.

[11] A simple modification of Theorem 1 allows us to construct equilibria with such features. Thus, such equilibria exist but are, arguably, of limited analytic interest.

Much added insight can be drawn by computing the ex ante equilibrium profit of a typical trader (see Corollary A.1). If we define $\mu(n) \equiv 1-(1-p)^n$, then the ex ante expected profit of a row-r trader is

$$\pi_r(\mathbf{n}) = \underbrace{\prod_{k=r+1}^{R} \mu(n_k)}_{[1]} \cdot \underbrace{p}_{[2]} \cdot \underbrace{(1-p)^{n_r-1}}_{[3]} \cdot \underbrace{\prod_{k=1}^{r-1} \delta(n_k) v}_{[4]} . \tag{2}$$

Expression (2) succinctly captures the complementary and competitive effects we mentioned above earlier. Complementaries flow from two sources.

- Term [1] captures the positive externality experienced by a row-r trader from an increase in the number of traders at upstream positions in the network. A trader earns profits only if the asset reaches his row and he is fortunate enough to buy and resell it. With increased upstream competition, this event becomes more likely as the risk of premature market breakdown recedes. $\mu(n)$ is the probability that at least one trader out of n has low trading costs. One low-cost trader is sufficient to ensure that trade does not break down at a particular row.

- Term [4] captures the positive externality experienced by a row-r trader from an increase in the number of traders at downstream positions in the network. It equals the asset's expected resale value. Thus, it summarizes the benefit from increased downstream competition, which inflates expected resale prices.

Terms [1] and [4], and therefore $\pi_r(n_r, \mathbf{n}_{-r})$, are increasing in $\mathbf{n}_{-r}$.

The direct competition that a trader experiences from others in the same row is captured by term [3]. It equals the probability with which a trader will be able to purchase the asset and resell it at a strictly positive profit. Term [3], and therefore $\pi_r(n_r, \mathbf{n}_{-r})$, is decreasing in n_r. Term [2] is simply the probably that a trader has low trading costs.

Remark 1. In (2), $\pi_r(\mathbf{n}) \propto v$. Thus, we henceforth normalize $v = 1$. (The normalization also applies to (3), defined below.)

3 Mergers and Stability

Whenever trade occurs through a network of intermediaries, two complementary questions arise. (1) Why does the network of intermediaries persist? And, (2) how did it arise? In this section, we focus on first of these questions. In section 4 we consider network formation.

In a stable network, existing market participants are willing to maintain the prevailing web of arms-length trading relationships. In practice, however, competition and complementarity may encourage agents to fold-in previously independent operations under a common umbrella. Firms often merge to constrain competition or to capitalize on complementary aspects of their operations. The former boosts market power while the latter expands scope. If arms-length economic relationships are to persist, integrative impulses must be kept at bay. Our definition of stability, proposed below, focuses precisely on such cases. We will call a network *stable* if collections of neighboring traders cannot profitably merge while performing the same intermediary task(s). Though intuitively simple, we work toward this definition by first outlining our model of mergers, which we call *partnership formation*.

A *partnership* is any connected subset of traders who merge and function as a single economic entity. Agents can form a partnership before private trading costs are realized but with knowledge of $\mathbf{n}$. We denote a partnership by a vector $\mathbf{m} = (m_1, \ldots, m_R)$ summarizing its composition. m_r is the number of traders from row r in the partnership $\mathbf{m}$.[12] As notation, we let $\bar{m} = \max\{r\colon m_r \geq 1\}$ and $\underline{m} = \min\{r\colon m_r \geq 1\}$ refer to the extreme rows occupied by members of $\mathbf{m}$. Traders not in a partnership are *independent traders*. Independent traders in rows $\underline{m}, \ldots, \bar{m}$ are said to be adjacent to the partnership $\mathbf{m}$.

The formation of a partnership changes the economy's structure. We assume that a partnership maintains all constituents' links to the wider economy, but it functions as a single actor thereby spanning multiple steps in the intermediation process. For example, Figure 3 shows the creation of a partnership $\mathbf{m} = (0, 2, 1, 0)$ in the network $\mathbf{n} = (4, 4, 3, 2)$. The partnership combines two row-2 traders with one row-3 trader. It has links to traders in rows 1 through 4.

Once established, a partnership can trade like a typical trader. It can buy and resell the asset via the prevailing protocol. It too incurs a trading cost, $\theta_{\mathbf{m}} \in \{0, t\}$. Generally, $p_{\mathbf{m}} = \Pr[\theta_{\mathbf{m}} = 0]$ will be a function of the partnership's composition, $\mathbf{m}$. To focus our analysis, however, we assume that

$$p_{\mathbf{m}} = \prod_{k=\underline{m}}^{\bar{m}} \mu(m_k) \tag{A-1}$$

for all $\mathbf{m}$. The motivation behind (A-1) is simple. If each trader's individual cost is low with probability p, then $\prod_{k=\underline{m}}^{\bar{m}} \mu(m_k)$ is the probability that there is at least one low-cost

[12]All agents in a particular row have the same neighbors. Therefore, this representation is without loss of generality given that we consider economies where there is at most one active partnership.

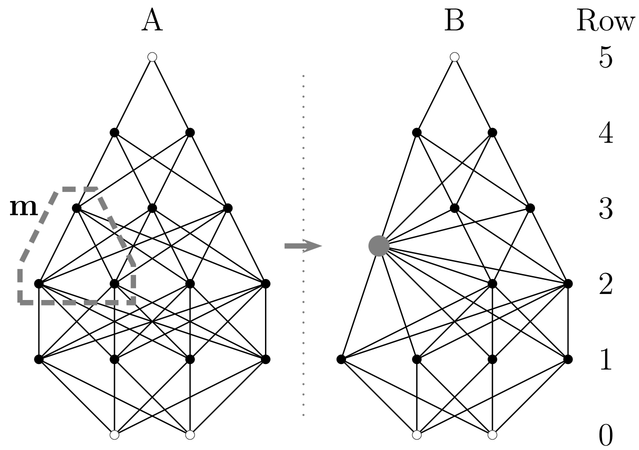

Figure 3: Formation of the partnership $\mathbf{m} = (0, 2, 1, 0)$. (Within-row links are omitted for clarity.)

Table 1: Initial equilibrium-path bids of low-cost agents in the networks of Figure 3.

Row	Network A	Network B	
		Independent Traders	Partnership
4	0.236	0.075	-
3	0.473	0.172	0.686
2	0.686	0.686* / ℓ**	
1	1	1	-

* When the asset is being sold by a row-3 trader.
** When the asset is being sold by the partnership.

trader in the partnership **m** from each row spanned by the partnership. An application for this specification could be a supply chain network where a sequence of distinct tasks must be performed and agents from different rows are specialized in those tasks.

As apparent from Figure 5, a partnership alters a network's structure. Though not immediate, the equilibrium of Theorem 1 has a natural generalization to the case a network with an active partnership. On the equilibrium path, traders and the partnership bid their expected resale values and the asset moves toward the buyers. Due to length, we formalize this equilibrium in Appendix B. Below we highlight its key novelties through an example. A partnership has both direct and indirect implications for the market's operation.

Example 2. Consider the two networks in Figure 3 and suppose $p = 1/2$. Theorem 1 describes bidding in Network A. High-cost traders always bid ℓ. Low-cost traders' equilibrium-path bids are defined inductively and are summarized in Table 1.[13]

Now consider Network B. Again, high-cost agents always bid ℓ. To characterize bidding by low-cost agents, we work up the rows of the network.

Row 1 On the equilibrium path, the problem faced by traders in row 1 is essentially the same as in Network A. Such traders bid 1.

Row 2 The optimal bid of a row-2 trader depends on the asset's seller.

1. Suppose the partnership is selling the asset. The asset's (unconditional) expected resale value to a row-2 independent trader is 0.686, as in Network A. However, a moment of reflection suggests this is an unwise bid for a row-2 trader. Given that row-1 agents are also neighbors of the partnership and are also bidding in the same auction, the asset's expected resale value to a row-2 trader *conditional* on winning this auction with a bid of 0.686 is zero. Given equilibrium play, a row-2 trader can acquire the asset from the partnership with a bid (strictly) less than 1 only when all traders in row 1 bid ℓ. However, this implies all traders in row 1 have high trading cost and a row-2 trader would not be able to profitably resell the asset. Hence, ℓ is an optimal bid.

2. Suppose a row-3 independent trader is selling the asset. If this event occurs on the equilibrium path, row-1 traders have not placed any bids and no value-relevant information is revealed during this sale. Hence, a row-2 trader can confidently bid his expected resale value, 0.686, in this contingency.

The Partnership The partnership's first opportunity to acquire the asset occurs when it is sold by a row-4 trader. In this case, a low-cost partnership can bid 0.686 as it can resell the asset at that expected price to a row-1 trader. It can be shown that a partnership cannot gain by instead waiting to purchase the asset from a trader in row 3 or row 2. Any possible benefits a delay may bring are already folded into the price it would pay conditional on acquiring the asset from a row-4 trader. (This price depends on the bids of row 3 traders, and so on.)

Row 3 The winner's curse intuition suggested in the case of a row-2 trader applies again to traders in row 3. Though a row-3 trader bids directly against the partnership when

[13] All values in this example are rounded to three decimal places.

the asset is sold by a row-4 trader, an agent in row 4 must anticipate reselling the asset for an expected price less than 0.686. His neighbors will not bid more than 0.686 and there is a chance some have high costs precluding resale altogether. If this trader acquires the asset with a bid less than 0.686, in equilibrium he ought to infer that the partnership has a high trading cost. Thus, he should adjust his bid accordingly to avoid a winner's curse. The asset becomes comparatively less valuable as this event signals reduced downstream competition. In equilibrium, he anticipates reselling it only to a low-cost row-2 trader. Thus, the asset's resale value to a row-3 trader is only $0.25 \times 0.686 = 0.172$.

Row 4 A row-4 trader may sell the asset to either an independent trader or to the partnership. With probability 0.281, the partnership has low costs and there is at least one low-cost independent trader. With probability 0.156 the partnership has high costs but both independent traders have low costs. In each case the sale price is 0.172. Hence, the asset's expected resale value is $(0.281 + 0.156) \times 0.172 \approx 0.075$, which defines an optimal bid.

Theorem 1's generalization to the case of a partnership builds on the preceding example's intuition. The key modification concerns the adjustment of independent traders' expected resale values to account for the "bad news" revealed when a multi-row partnership fails to acquire the asset upon its first chance. Such inferences are important as the same agent may participate in multiple auctions on the equilibrium path thereby revealing information about their trading costs. We can gleam further insight into this effect by decomposing the partnership's ex ante equilibrium profit. As shown in Appendix B,

$$\pi_{\mathbf{m}}(\mathbf{n}) = \underbrace{\prod_{k=\bar{m}+1}^{R} \mu(n_k)}_{[1]} \cdot \underbrace{\prod_{k=\underline{m}}^{\bar{m}} \mu(m_k)}_{[2]} \cdot \underbrace{\left(1 - \overbrace{\mu(n_{\bar{m}} - m_{\bar{m}})}^{[3a]} \overbrace{\prod_{k=\underline{m}}^{\bar{m}-1} \delta(n_k - m_k)}^{[3b]}\right)}_{[3]} \cdot \underbrace{\prod_{k=1}^{\underline{m}-1} \delta(n_k)}_{[4]}. \quad (3)$$

The labeling of (3) parallels that of (2) for the baseline model.[14] Term [1] captures the benefit from increased upstream competition while term [4] is the asset's expected resale value given the normalization $v = 1$. Term [2] is the probability with which the partnership has low trading costs given (A-1). This term is increasing in m_k but decreasing in $(\bar{m} - \underline{m})$ and it summarizes the partnership's trading technology. An increase in $\bar{m} - \underline{m}$ corresponds to an expansion of the partnership's scope as it moves into additional intermediary tasks. Term

[14]If a partnership has one member, (3) collapses to (2) given the convention $\prod_{k=r}^{r-1} \delta(n_k - m_k) = 1$.

Table 2: Benefits and Costs of Exclusively Vertical and Exclusively Horizontal Mergers

	Terms in (3)	Horizontal Mergers	Vertical Mergers
Trading Technology	[2]	+	-
Direct Market Power	[3a]	+	
Indirect Market Power	[3b]		+
Distance Premium	[1], [4]		+

[3] accounts for the partnership's market power in the network. It has two key elements. Term [3a] captures the direct decline in competition due to the partnership's presence. The partnership bids against fewer competitors and thus it can secure more favorable terms more often. Term [3b] captures an indirect market power enhancement flowing from the partnership's informational advantage. Due to its scope, a partnership has better knowledge concerning the intensity of downstream competition in comparison to independent traders. Specifically, the partnership knows whether it is a potential purchaser of the asset from downstream traders. In response to this informational advantage, independent traders must temper their bids to avoid the winner's curse effect noted above and illustrated in Example 2. The reduction of independent traders' resale values propagates through the network and deflates the bids of independent traders in row $\bar{m}$, further reducing the partnership's expected payment conditional on acquiring the asset.

As suggested by the above discussion, merging along vertical and horizontal dimensions can have different implications. These are summarized in Table 2. A purely horizontal merger improves the partnership's trading technology and gives the partnership direct market power. Unambiguously, these enhance profits. Vertical integration implies an expanded scope as the partnership spans multiple steps in the intermediation chain. The need to accomplish multiple intermediation tasks weakens the group's technology given (A-1). On the other hand, a vertical partnership enjoys some indirect market power and a direct distance premium. This final effect is purely mechanical. The partnership buys at a low price from an agent in row $\bar{m}+1$ and sells at a premium to an agent in row $\underline{m}-1$. Partnerships that combine vertical and horizontal links, like in Figure 3, experience some mixture of these benefits and costs.

Though we have already identified an indirect cost of merging, we have yet to consider the direct costs that mergers entail in practice. For example, it is often costly to integrate the operations and cultures of two previously separate firms. Legal constraints, such as antitrust laws, can make collusive arrangements or mergers difficult. To capture direct costs we further assume that when a partnership $\mathbf{m}$ forms, it incurs a cost of $\zeta(\mathbf{m})$. Though mindful of more general specifications, for simplicity we assume that

$$\zeta(\mathbf{m}) = \underbrace{c_h \sum_{r=\underline{m}}^{\bar{m}} (m_r - 1)}_{[1]} + \underbrace{c_v \cdot (\bar{m} - \underline{m})}_{[2]} \tag{A-2}$$

where $c_h, c_v \geq 0$ are constants. (A-2) distinguishes between two kinds of merging actions.[15] Term [1] captures the cost of fusing *horizontal* links in the network. Term [2] captures the costs of fusing *vertical* links in the network.

3.1 Network Stability

We call a network stable if no partnership can provide its members a greater payoff relative to a benchmark where all agents act independently.

Definition 1. A trading network $\mathbf{n}$ is *stable* if for all feasible partnerships $\mathbf{m} = (m_1, \ldots, m_R) \leq \mathbf{n}$, $\sum_r m_r \pi_r(\mathbf{n}) \geq \pi_{\mathbf{m}}(\mathbf{n}) - \zeta(\mathbf{m})$.

Our definition of stability draws inspiration from classic solution concepts, such as the core, in a transferable-utility setting. By focusing on the fusing of nodes it contrasts with other common definitions of stability in network economies. For example, Jackson and Wolinsky (1996) propose a stability notion whereby a fixed set of agents can form and drop links. Ostrovsky (2008) models supply-chain networks and proposes a generalization of "stability," in the sense of Gale and Shapley (1962), to that class of problems. Manea (2014) examines horizontal and vertical integration of traders in a network like we do, but focuses on comparative static welfare implications rather than network stability.

Whether a trading network is stable is closely related to the underlying trading technology and the magnitude of merger costs.

Theorem 2. *If $c_h > 0$ and $c_v \geq 0$, then there exists a $\hat{p} > 0$ such that for all $p < \hat{p}$, the trading network is stable.*

Theorem 2 shows that a stable network exists when p is sufficiently small. In this case, traders frequently experience costs shocks. The insulation provided by a web of independent

[15](A-2) may be viewed as a first-order approximation to a more general—for example, convex—cost function.

trading partners is particularly valuable in this case and acts as a natural disincentive to integrative actions. Curiously, the drawbacks of vertical partnership formation (see Table 2) may be sufficiently strong so that stability can be assured even if direct, vertical merger costs are zero.[16] In contrast, if $c_h = 0$, then instability is virtually assured.

Theorem 3. *If $n_1 \geq 2$ and $c_h = 0$, then the trading network is not stable.*

Theorem 3 highlights the differential impact of horizontal and vertical merger costs. Notably, a network can be stable even if vertical merger costs are zero. This is the outcome, for example when vertical mergers are associated with a pronounced deterioration in trading technology and enhanced scope is not profit enhancing. Purely horizontal mergers enhance traders' market power and improve their trading technology. Therefore, some direct costs must counteract these benefits to ensure stability.

3.2 Instability and Welfare

While all trading networks are stable if direct merger costs are sufficiently large, instability may ensue if such costs are small. If a network is not stable, what might happen? Perhaps the simplest consequence is that the economy operates with a partnership in its midst, at least in the short-run. Although this arrangement may not persist in the longterm, it provides an obvious benchmark to gauge the welfare implications of this alternative market structure. Intuitively, one may interpret a market with an active, multi-row partnership as a market that is in the initial phases of "disintermediation." The minimal economic distance between the buyers and the seller is shorter than it was initially.

To measure welfare in our economy, we first define

$$\chi(\mathbf{n}) = \prod_{r=1}^{R} \mu(n_r)$$

as the market's *capacity*. It is the probability that the asset reaches a buyer given the configuration $\mathbf{n}$. Therefore, it accords naturally with the market's throughput. Of course, $\chi(\mathbf{n})$ also equals the expected surplus generated in the economy. Therefore, it provides a meaningful, utilitarian welfare measure.

[16] This observation accommodates Ostrovsky's (2008, p. 911) argument that in a trading network it may be easier to organize a "vertical" coalition than a horizontal one. A purely vertical partnership shares a structure with the "chain block" proposed by Ostrovsky (2008).

Theorem 4. *A network's capacity, $\chi(\mathbf{n})$, equals the sum of the intermediary traders' expected profits, the expected profit of the seller, and the expected payoff of the buyers.*

If there is an active partnership $\mathbf{m}$, the market's capacity becomes

$$\chi_{\mathbf{m}}(\mathbf{n}) = \prod_{r=\bar{m}+1}^{R} \mu(n_r) \left[\prod_{r=\underline{m}}^{\bar{m}} \mu(m_r) + \left(1 - \prod_{r=\underline{m}}^{\bar{m}} \mu(m_r)\right) \prod_{r=\underline{m}}^{\bar{m}} \mu(n_r - m_r) \right] \prod_{r=1}^{\underline{m}-1} \mu(n_r).$$

By inspection, two conclusions are immediate.[17] First, if a partnership is confined entirely to a single row, its presence does not impact aggregate welfare: $\chi_{\mathbf{m}}(\mathbf{n}) = \chi(\mathbf{n})$. Though the partnership has enhanced market power, it only introduces distributional consequences with no impact on the market's efficiency. Second, and contrasting the first observation, if $\mathbf{m}$ spans multiple rows, then $\chi_{\mathbf{m}}(\mathbf{n}) < \chi(\mathbf{n})$. That is, the partnership's presence not only alters the distribution of benefits among traders, but it also reduces aggregate welfare—a deadweight loss not unlike in the case of a classic monopoly.

The preceding discussion complements Theorems 2 and 3 and links their conclusions with aggregate welfare. Notably it suggests that idiosyncratic risk, here modeled as cost shocks, serves to reinforce a relatively more efficient market organization. Even if direct merger costs are small, the relatively more efficient market configuration can be maintained as multi-row partnerships are financially unrewarding.

4 Entry and Equilibrium Networks

While stability concerns the persistence of an existing network of trading relationships, it does not address the process governing network formation. We assume that the network formation process is characterized by the free entry of intermediaries given a fixed entry cost. Though distinct from most models of network formation,[18] our model shares numerous features with many classic models in industrial organization or international trade theory.

Fix R and suppose there is a large group of potential traders who may enter the market at any of the R levels while forming links to agents in adjacent positions. To enter the market, a trader must incur an entry cost of $\kappa > 0$. We interpret κ as an irreversible investment in market-specific skills or technology. For example, it may be the cost of forming relevant relationships to be a part of the trading community. Once all traders have made their entry

[17]Echoing Theorem 4, $\chi_{\mathbf{m}}(\mathbf{n})$ can be shown to equal a sum of expected profits.
[18]See Jackson (2008, Part 2) for a survey of network formation.

and location decisions, the network configuration becomes known, traders learn their costs, and exchange unfolds as before. Agents not entering the market receive a payoff of zero. Entry occurs until no further profitable entry is possible.

Definition 2. The network configuration $\mathbf{n}^* = (n_1^*, \ldots, n_R^*)$ is an *equilibrium configuration* if for all r, $\pi_r(\mathbf{n}^*) - \kappa \geq 0$ and $\pi_r(n_1^*, \ldots, n_{r-1}^*, n_r^* + 1, n_{r+1}^*, \ldots, n_R^*) - \kappa < 0$.

Definition 2 translates the standard intuition associated with free entry, i.e. profits being driven to zero, to our setting. Our definition is closely related to the "equilibrium configurations" analyzed by Gary-Bobo (1990) in a class of asymmetric entry models. Our study is outside that paper's purview since traders' payoffs in our model do not satisfy his monotonicity condition.

4.1 The Set of Equilibrium Networks

All markets feature an equilibrium configuration. This conclusion is immediate when $R = 1$. When $R \geq 2$ there exists a trivial equilibrium with no traders.[19] Although an important case—speculatively, many unobserved markets do not exist because of "coordination" on the no trade equilibrium—this equilibrium is of limited analytic interest. More interestingly, however, nontrivial equilibria exist under mild conditions.

Theorem 5. *Let* $\bar{n} \equiv \left\lceil 1 + \frac{\log(\kappa) - \log(p)}{\log(1-p)} \right\rceil$ *and define* $\bar{\mathbf{n}} = (\bar{n}, \ldots, \bar{n})$.

1. *If* $\mathbf{n}^*$ *is an equilibrium, then* $\mathbf{n}^* \leq \bar{\mathbf{n}}$.[20]
2. *There exists a nontrivial equilibrium if and only if there exists some* $\mathbf{n} \leq \bar{\mathbf{n}}$ *such that* $\pi_r(\mathbf{n}) - \kappa \geq 0$ *for all* r.

The proof of Theorem 5 defines a tâtonnement process that monotonically converges to an equilibrium. The process begins from an initial configuration where agents in each row earn positive expected profits. The number of traders in each row is then increased successively until the profits of a typical trader satisfy the conditions of Definition 2.

Though equilibria exist, the presence of complementarities implies that they are often not unique. Surprisingly, however, the set of equilibria has a particularly tractable structure allowing for meaningful comparisons and welfare analysis. It is a directed set and one

[19] Conditional on an empty network, entry by a single agent is always unprofitable. Either the agent cannot acquire the asset or he has no one to sell it to.

[20] We employ the usual coordinate-wise partial ordering of vectors.

equilibrium—the maximal equilibrium—dominates others in terms of the intensity of trader competition. First, we illustrate these conclusions with an example. Thereafter we formalize them in Theorem 6.

Example 3. Suppose $R = 6$, $p = 0.5$, and $\kappa = 0.01$. Given these parameters, there exist two equilibrium networks: $\mathbf{n} = (4, 4, 3, 3, 2, 1)$ and $\mathbf{n}' = (6, 6, 6, 6, 5, 5)$. See figures 4 and 5. (Like in all diagrams to follow, we omit within-row links for clarity.) Clearly, network $\mathbf{n}'$ is the maximal equilibrium given this parameterization.

Theorem 6. *Let* $\mathbf{n}$ *and* $\mathbf{n}'$ *be equilibria. There exists an equilibrium* $\mathbf{x}$ *such that* $\mathbf{x} \geq \mathbf{n}$ *and* $\mathbf{x} \geq \mathbf{n}'$*.*

To prove Theorem 6 we rely on the same tâtonnement process as in the proof of Theorem 5; however, $\mathbf{n} \vee \mathbf{n}' = (\max\{n_1, n_1'\}, \ldots, \max\{n_R, n_R'\})$ serves as the initial condition. As all equilibrium networks are bounded above by $\bar{\mathbf{n}}$, successive applications of Theorem 6 lead to the following corollary.

Corollary 1. *There exists an equilibrium* $\mathbf{q}^*$ *such that* $\mathbf{q}^* \geq \mathbf{n}^*$ *for every other equilibrium* $\mathbf{n}^*$*. We call* $\mathbf{q}^*$ *the maximal equilibrium.*

Among all equilibria, the maximal equilibrium features the most intensive competition among traders. In every row the maximal equilibrium has the most traders. As we explain below, this fact has important welfare and market-robustness implications.

4.2 Equilibrium Configurations and the "Bullwhip Effect"

Like the equilibrium set, individual equilibria also have a predictable and tractable structure. Consider again Example 3. Though more visible in $\mathbf{n}$ than in $\mathbf{n}'$, both networks share a pyramid-like form. There are more intermediary traders near the buyers than the seller. This is a characteristic of all equilibrium markets.

Theorem 7. *If* $\mathbf{n}^* = (n_1^*, \ldots, n_R^*)$ *is an equilibrium network, then* $n_r^* \geq n_{r+1}^*$*.*

The logic behind Theorem 7 is easily illustrated with a thought experiment. Suppose a market has an equal number of traders in each row; that is, the market is balanced in its distribution of traders. Given this market configuration, however, there is an imbalance in the profits of traders in different rows. Since $\mu(n_r) > \delta(n_r)$, the expected profits of a row-1 trader are greater than the expected profits of row-R trader when $n_1 = n_R$. If entry costs are

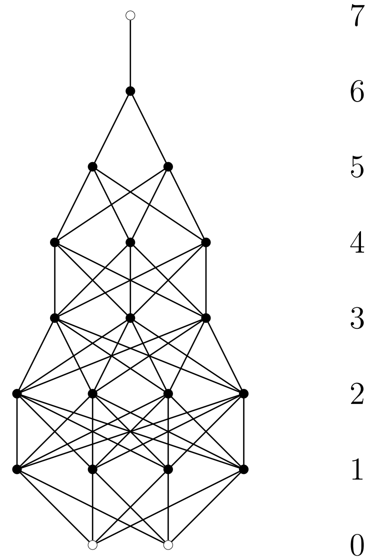

Figure 4: The equilibrium $\mathbf{n}$ in Example 3. (Within-row links are omitted for clarity.)

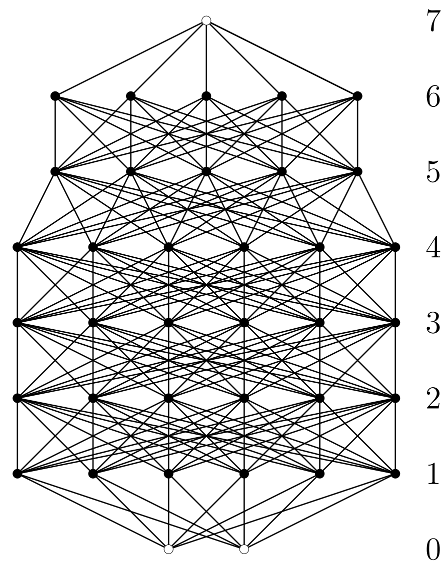

Figure 5: The equilibrium $\mathbf{n}'$ in Example 3. (Within-row links are omitted for clarity.)

sufficiently low, additional traders would be attracted to positions near the buyer thereby generating the skewed distribution of intermediaries.

To add economic rationale to the preceding explanation, it is helpful to focus on the different types of uncertainty encountered by different traders. $\mu(n_r)$ equals the probability that there is at least one low-cost trader in row r. Thus, it is the probability that the asset successfully transits a level of the network. $\delta(n_r)$, on the other hand, is the expected fraction of the resale value that an agent can appropriate from a sale. A trader earns profits if there are at least two low-cost agents who are potential buyers. Thus, $\delta(n_r)$ captures additional uncertainty concerning the terms of trade governing subsequent transactions. With respect to upstream transactions, a trader only cares that they occur. With regards to downstream transactions, a trader also cares about the prices at which they occur. By moving closer to the buyer, aggregate price uncertainty diminishes as there are few downstream transactions. Furthermore, ever-enhancing competition additionally reduces price variability.

The Bullwhip Effect As noted above, one possible interpretation of our model is that of a supply chain. A stylized fact observed in many supply chains is the "bullwhip effect" (Lee et al., 1997a,b).[21] Roughly, this effect corresponds to an increase in the variability of demand at higher levels in a supply chain. Though distinct from the four explanations proposed by Lee et al. (1997b), our model is nevertheless consistent with this stylized fact when we consider equilibrium network configurations. For example, if we focus on the sale of the asset by an agent in row $r+1$ to an agent in row r, we can define the coefficient of variation in demand, $CVD_r(\mathbf{n})$, as

$$CVD_r(\mathbf{n}) = \frac{\text{Standard Deviation of Demand}}{\text{Expected Demand}} = \sqrt{\frac{1}{\mu(n_r)} - 1}.$$

Similarly, we can define the coefficient to variation in sales price, $CVP_r(\mathbf{n})$, as

$$CVP_r(\mathbf{n}) = \frac{\text{Standard Deviation of Price}}{\text{Expected Price}} = \sqrt{\frac{1}{\delta(n_r)} - 1}.$$

We detail the derivation of these terms in Appendix C.

Since $\mu(n_r)$ and $\delta(n_r)$ are increasing in n_r, $CVD_r(\mathbf{n})$ and $CVP_r(\mathbf{n})$ are decreasing in n_r. In an equilibrium network, $n_r^* \geq n_{r+1}^*$; thus, relative variation in demand and prices increases

[21]We are grateful to Vasco Carvalho for bringing this phenomenon to our attention.

r	1	2	3	4	5	6
$CVD_r(\mathbf{n})$	0.258	0.258	0.378	0.378	0.577	1
$CVD_r(\mathbf{n}')$	0.126	0.126	0.126	0.126	0.180	0.180
$CVP_r(\mathbf{n})$	0.674	0.674	1	1	–	–
$CVP_r(\mathbf{n}')$	0.350	0.350	0.350	0.350	0.480	0.480

Table 3: Coefficient of variation in demand and price in Example 3.

as one moves away from consumers. Moreover, $CVP_r(\mathbf{n}) \geq CVD_r(\mathbf{n})$, which implies greater relative variation in prices than in demand. Table 3 provides a sense of the magnitudes of these values in the equilibrium networks of Example 3.

4.3 Equilibrium and Welfare

Above we defined a market's capacity, $\chi(\mathbf{n})$, as the probability the asset traverses the network. Theorem 4 showed that $\chi(\mathbf{n})$ equals the sum of agents' expected profits. By inspection, we can conclude it is increasing in $\mathbf{n}$. Therefore, it is clear that maximal equilibria enjoy a welfare advantage under this metric. For instance, in Example 3 the sparse equilibrium network $\mathbf{n}$ has a capacity of approximately 0.25. The maximal equilibrium network's capacity is 0.88. The sparse network's capacity is particularly impacted by its characteristic bottleneck around row 6. When a cost shock hits a trader in row 6, or even in row 5, its effect is amplified as few other agents can function as effective substitutes for traders in those positions. The result is that welfare is compromised.

While capacity may be an appropriate welfare metric for an existing network, from an ex ante point of view it may be inadequate as it ignores incurred entry costs. Pursuing this vein, we define

$$\Omega(\mathbf{n}) = \chi(\mathbf{n}) - \kappa \sum_{r=1}^{R} n_r \tag{4}$$

as the (ex ante) aggregate welfare generated by a network. The welfare-dominance of the maximal equilibrium under this metric is no longer obvious from inspection. The network's capacity increases in $\mathbf{n}$, but does aggregate entry cost. Rewriting (4) as the sum of buyers', traders', and the seller's payoffs yields a helpful decomposition:

$$\Omega(\mathbf{n}) = \underbrace{n_0\pi_0(\mathbf{n})}_{\text{Buyers' Payoffs}} + \underbrace{\sum_{r=1}^{R} n_r(\pi_r(\mathbf{n}) - \kappa)}_{\text{Traders' Payoffs}} + \underbrace{\pi_{R+1}(\mathbf{n})}_{\text{Seller's Payoff}} .$$

If $\mathbf{n}^*$ is an equilibrium configuration, then $\pi_0(\mathbf{n}^*) = 0$. Moreover, $\pi_r(\mathbf{n}^*) - \kappa \approx 0$ for all $1 \leq r \leq R$ due to free entry.[22] Thus, in an equilibrium configuration

$$\Omega(\mathbf{n}^*) \approx \pi_{R+1}(\mathbf{n}^*) = \prod_{r=1}^{R} \delta(n_r^*),$$

which is increasing in $\mathbf{n}^*$. Therefore, aggregate welfare increases with the number of equilibrium traders and the maximal equilibrium remains favored.

While the maximal equilibrium configuration offers compelling welfare advantages, it falls short of the welfare-maximizing configuration. When a trader enters the market, he imparts a positive externality on traders located at other levels of the network, boosting their profits. Since traders do not internalize this benefit, under-entry relative to a first-best benchmark is a possible outcome. A countervailing force exists, however, as a trader's entry imparts a negative externality on his direct competitors who co-locate at the same level. The pursuit of profit, analogous to "business stealing," may encourage an over-entry of intermediaries. The following theorem confirms that the former effect dominates.

Theorem 8. *Let $\hat{\mathbf{n}}$ be the ex ante welfare maximizing network configuration. That is, $\hat{\mathbf{n}}$ solves*

$$\max_{\mathbf{n}} \ \Omega(\mathbf{n}). \tag{OPT}$$

For all r and r', $\hat{n}_r = \hat{n}_{r'}$. Moreover, if $\mathbf{n}^*$ is an equilibrium configuration, then $\hat{\mathbf{n}} \geq \mathbf{n}^*$.[23]

The first part of Theorem 8 concludes that a welfare-maximizing network equalizes the number of intermediaries across rows. It is a consequence of $\Omega(\mathbf{n})$'s symmetry, which is clear from (4). The second part follows from the presence of externalities. The wedge between the private profits motivating entry and the social benefits associated with a dense set of intermediaries leads to intermediary under-entry in equilibrium.

To reinforce ideas, we provide two examples highlighting the relationship between equilibrium networks and welfare.

Example 4. Suppose $R = 5$ and $p = 1/2$. By varying κ we can trace out a family of equilibria with differing welfare properties. The results of this experiment are summarized by Table 4. For each value of κ, the table presents all equilibrium configurations, $\mathbf{n}^* = (n_1^*, \ldots, n_5^*)$, along with the corresponding values for aggregate welfare, $\Omega(\mathbf{n}^*)$, and capacity, $\chi(\mathbf{n}^*)$. Analogous

[22]The integer constraint prevents exact equality.

[23]In the non-generic case where (OPT) has multiple solutions, we assume $\hat{\mathbf{n}}$ is the greatest solution.

Table 4: Equilibrium and welfare-maximizing (OPT) networks in Example 4.

κ	Equilibrium							OPT		
	n_1^*	n_2^*	n_3^*	n_4^*	n_5^*	$\Omega(\mathbf{n}^*)$	$\chi(\mathbf{n}^*)$	$\hat{n}$	$\Omega(\hat{\mathbf{n}})$	$\chi(\hat{\mathbf{n}})$
0.005	7	7	7	7	7	0.79	0.96	7	0.79	0.96
0.010	6	6	6	6	5	0.62	0.91	6	0.62	0.92
0.015	5	5	5	5	4	0.47	0.83	5	0.48	0.85
	4	3	3	2	1	0.07	0.27			
0.020	5	4	4	4	3	0.30	0.70	5	0.35	0.85
	4	4	3	3	2	0.18	0.50			
0.022	4	4	3	3	2	0.15	0.50	5	0.33	0.85
0.025	0	0	0	0	0	0	0	5	0.23	0.85
0.030	0	0	0	0	0	0	0	4	0.12	0.72

values for the welfare-maximizing network, $\hat{\mathbf{n}} = (\hat{n}, \ldots, \hat{n})$, are also provided. When $\kappa = 0.005$, $\mathbf{n}^*$ coincides with $\hat{\mathbf{n}}$. As κ increases, we observe both equilibrium multiplicity and a divergence between $\Omega(\mathbf{n}^*)$ and $\Omega(\hat{\mathbf{n}})$. As expected, more imbalanced equilibria imply a greater welfare loss. When κ is sufficiently large the market fails to operate even though a socially-optimal configuration could generate a positive aggregate surplus.

Example 5. Suppose $p = 1/3$ and consider network lengths of $R \in \{4, 5, 6\}$. For each network, Figure 6 presents the aggregate welfare, $\Omega(\mathbf{n}^*)$, associated with all equilibria as a function of κ. The dashed curves indicate the corresponding first-best welfare levels. When κ is low, "thick" equilibria prevail and the planner's solution aligns closely with the unique equilibrium. The welfare gap between first-best and equilibrium networks widens as κ increases. Small changes in κ often shift the economy between equilibria.

The above discussion provides at best a mixed conclusion concerning the welfare properties of equilibrium networks. The under-entry of intermediaries into the market begets a direct welfare loss relative to the first-best benchmark. Compounding that loss, however, is the specific configuration assumed by traders who do enter the market. A network can simultaneously function both as an absorber and as an amplifier of idiosyncratic risks.[24] The "pyramid" network structure exaggerates the latter. The market is disproportionately sensitive to the shocks experienced by the few traders located near the seller. Such agents have few close substitutes and provide important complementarity to downstream agents. This conclusion is in line with that of Acemoglu et al. (2012) who show that small sectoral

[24]We thank Richard Zeckhauser for suggesting to us the amplifier/absorber metaphor.

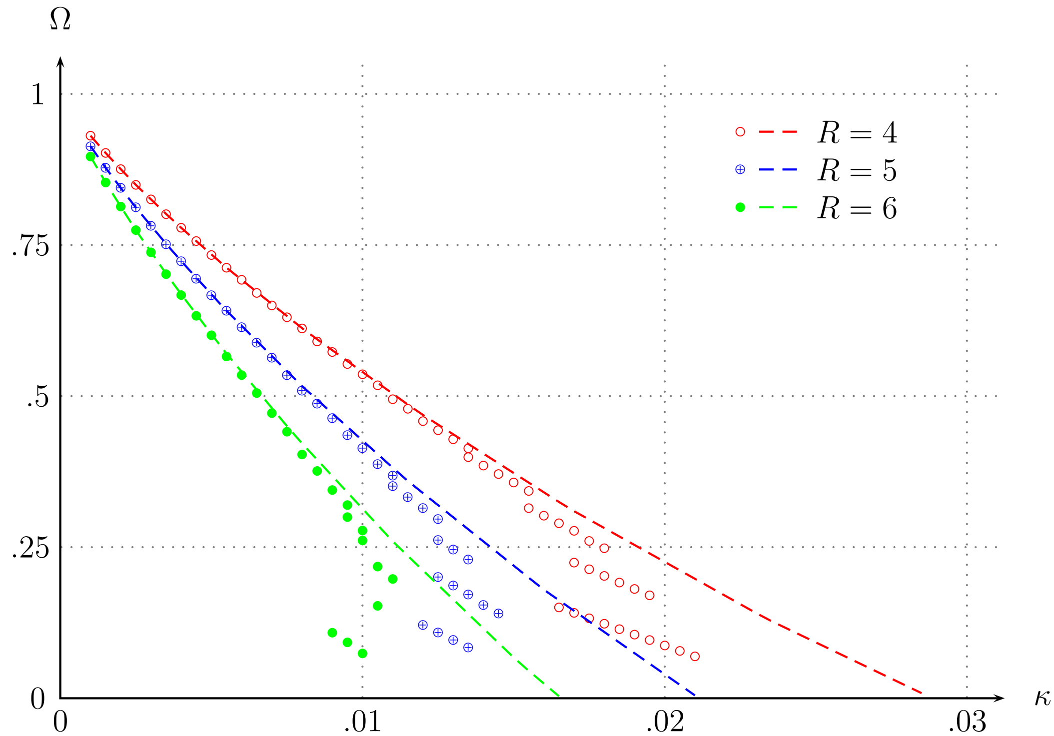

Figure 6: Welfare and entry costs in Example 5.

shocks can have a disproportionately-large impact in a macroeconomic context due to an economy's network structure.

Several policy tools are available to address the inefficiencies that we have identified. First, policies that remove entry barriers (i.e. decrease κ) are a safe bet in moving the economy toward a more efficient organization. This claim is hardly novel, but it is reinforces classic insights. Likewise, the presence of externalities admits the possibility of Pigouvian subsidies as a corrective policy. Notably, the seller, who is the main beneficiary of a more dense network, could subsidize the entry of intermediaries. Finally, a more subtle policy change may modify the institutions or protocols governing exchange, perhaps on a location-by-location basis. For instance, if traders in row R could impose a small reserve price, their profits would be enhanced. This may justify further entry into that location.[25] Institutional changes may not be costly to implement in a direct sense, but may be indirectly costly due to inertia in market culture or practice.

5 Equilibrium and Stability

In our terminology, equilibrium network configurations (consistent with free-entry) and stable network configurations (immunity to mergers) are independent concepts. However, they naturally work together. For example, one might ask which equilibrium configurations (if there are many) are more inclined to be stable?

For illustration, consider an economy where $R = 2$ and let $\mathbf{n}^* \leq \mathbf{n}^{**}$ be equilibrium networks. Although many partnerships may serve to destabilize this network, for brevity consider a partnership $\mathbf{m}$ that includes traders only from row $r \in \{1, 2\}$. This partnership's expected profit can be written as

$$\pi_{\mathbf{m}}(\mathbf{n}) = \pi_r(\mathbf{n}) \frac{1-p}{p} \left(\frac{1}{(1-p)^{m_r}} - 1 \right)$$

where

$$\pi_r(\mathbf{n}) = \begin{cases} \mu(n_2) p (1-p)^{n_1 - 1} & \text{if } r = 1 \\ p(1-p)^{n_2 - 1} \delta(n_1) & \text{if } r = 2 \end{cases}$$

is the expected profit of a typical row-r trader when $R = 2$.

[25]Crucially, however, the trading mechanism cannot be tilted too strongly to traders in row R. Else, they may extract too much surplus from traders in row $R - 1$ rendering entry in that location unattractive. Striking the right balance in terms of bargaining power would be crucial.

As in our discussion of welfare, economic comparisons of equilibrium networks are simplified due to the free-entry of intermediaries. When the economy is in equilibrium, expected trader profits are pinned-down by entry costs (with an allowance to account for the integer constraint). Thus, $\pi_r(\mathbf{n}^*) \approx \pi_r(\mathbf{n}^{**}) \approx \kappa$ and so $\pi_{\mathbf{m}}(\mathbf{n}^*) \approx \pi_{\mathbf{m}}(\mathbf{n}^{**})$.

For instance, suppose that when $m_r \leq n_r^*$,

$$\pi_{\mathbf{m}}(\mathbf{n}^{**}) - \zeta(\mathbf{m}) \leq m_r \pi_r(\mathbf{n}^{**}). \tag{5}$$

That is, the large equilibrium network cannot be destabilized by a relatively small partnership, which could also form in the smaller network. In this case, the same inequality should also obtain when the underlying network is $\mathbf{n}^*$. On the other hand, there might exist a feasible partnership $\mathbf{m}' \leq \mathbf{n}^{**}$ that may compromise the large market's stability but which is infeasible in the small market (i.e. $m'_r > n^*_r$). Since $\pi_{\mathbf{m}'}(\mathbf{n}^{**})$ is increasing in m'_r, for m'_r sufficiently large the inequality in (5) can reverse: $\pi_{\mathbf{m}'}(\mathbf{n}^{**}) - \zeta(\mathbf{m}') > m'_r \pi_r(\mathbf{n}^{**})$. Thus, the critical-mass of active traders in the large market, though beneficial from a welfare perspective, can be a risk factor pulling toward market instability. As illustrated by the following example, similar conclusions also appear in more complex economies.

Example 6. Suppose $R = 5$ and $p = 1/2$. When $\kappa = 0.015$, there are two equilibrium configurations: $\mathbf{n}^* = (4, 3, 3, 2, 1)$ and $\mathbf{n}^{**} = (5, 5, 5, 5, 4)$.[26] If merger costs conform to (A-2), both networks are stable when c_h and c_v are large. Stability is compromised when c_h and c_v are low. Specifically, Figure 7 identifies the stable network(s) for each pair of parameters (c_h, c_v). The network $\mathbf{n}^{**}$ is less robust than $\mathbf{n}^*$ as it is stable only when direct merger costs are greater.

While we have already emphasized the welfare benefits of maximal equilibrium configurations, the above discussion suggests ensuring those benefits may be difficult. First, to arrive upon a maximal equilibrium, agents' entry decisions must be coordinated to leverage the benefits of complementarities in distant regions of the network. As is well-known, coordination on the "good equilibrium" is never assured. Second, even if entry challenges can be overcome, ensuring market stability may prove more challenging than had a non-maximal equilibrium configuration prevailed. If a partnership forms, for example, the expected welfare gains of a maximal equilibrium do not materialize fully. When stability is an important consideration or constraint, a non-maximal equilibrium configuration may be the best feasible outcome.

[26] This case was examined in Example 5 as well.

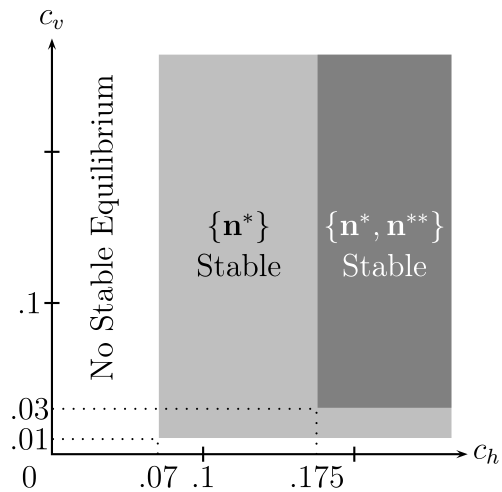

Figure 7: Stable equilibria as a function of (c_h, c_v) in Example 6.

6 Context, Extensions, and Conclusions

We have developed a model of network formation highlighting the competition and complementarity among intermediaries. These forces shape both network formation and affect the persistence or stability of existing networks. Our model shows that markets may not naturally assume the most capable market organization. The bipartite buyer-seller networks traditionally explored in the literature do not always identify these effects as the complementarities among agents are tempered by the assumed network structure.

Demand vs. Supply Uncertainty A key ingredient fueling many of our results is that intermediaries face *demand* uncertainty. Traders hold residual uncertainty regarding the asset's liquidity as their neighbors are exposed to private cost shocks, which may preclude exchange. Asymmetries in the nature of uncertainty lent equilibrium networks their characteristic, pyramid-like structure with more traders congregating near the buyers than the seller.

While demand uncertainty is present in many markets with active intermediaries, some markets—such as those for some commodities—feature *supply* uncertainty. Playing our model "in reverse" provides a framework for analyzing markets operating within this paradigm. Briefly, such a market could function as follows. A single buyer (in row 0) wishes to acquire an asset, which is supplied by multiple sellers (in row $R + 1$). To purchase this asset, the buyer contacts his neighboring intermediaries (in row 1). The buyer holds a procurement auction and the intermediary offering the lowest price is contracted to supply the asset. This auction could be implemented as a descending clock auction, mirroring the ascending auction that could be used in our original model.[27] Intermediaries with low trading costs submit competitive bids while those with high trading costs bid ℓ. If all intermediaries bid ℓ, they cannot supply the asset and the transaction breaks down. If an intermediary wins the procurement auction, he must now secure supply of the asset. To do so, the intermediary (in row 1) himself organizes a procurement auction, which now draws neighboring intermediaries in row 2. The process repeats until an intermediary secures supply of the asset from a seller in row $R+1$. If the process does not break down, a chain of low-cost intermediaries will link the buyer to a seller thereby allowing for exchange.

It is clear that competition and complementarity operate in the "reversed" market in much the same way as they did in our original economy. Parallel conclusions follow. Intermediary under-entry relative to a socially-optimal benchmark and a regularized network structure—in this case a "funnel" instead of pyramid—continue to feature in equilibrium configurations. The definition of market stability translates verbatim to this setting as well. Therefore, our basic framework can be adapted to accommodate many alternative trading structures, with under-entry remaining the underlining theme.

Related Literature In studying intermediation, our study builds on earlier analyses in several literatures. Networks provide a natural forum for studying exchange and the relationships among economic agents. In particular, our equilibrium stresses the complementaries among agents in the presence of network externalities (Economides, 1996).[28] Intuitively, traders who perform similar tasks in the intermediation process (i.e. those who have the same "friends") function as substitutes. In contrast, traders who are in distant regions of the economy complement each other. Downstream traders enhance competition and thus bid up resale prices. Upstream traders enhance the frequency of exchange; idiosyncratic shocks are less likely to compromise the market's operation.

Like Bala and Goyal (2000), Kranton and Minehart (2001), or more recently Condorelli and Galeotti (2012b), we study network formation. Our network-formation process builds around free entry and contrasts with their focus on strategic link formation. Additionally, our analysis moves away from bipartite buyer-seller networks by incorporating layers of intermediaries or middlemen. In this regard, our study follows most closely recent work

[27] Alternatively, intermediaries may engage in Bertrand competition iteratively lowering their offer prices lower and lower.

[28] Jackson (2008) provides a comprehensive survey of the literature on economic networks.

by Gale and Kariv (2007, 2009) who also study intermediation with a network of successive intermediaries.[29] Unlike these papers we endow traders in our model with private information about trading costs. Recognizing the importance of market "middlemen," Rubinstein and Wolinsky (1987) offer a lucid analysis based on the random matching of buyers and sellers with intermediaries. They do not explicitly model a network but their model accommodates alternative institutional arrangements, such as consignment sales, which we do not consider.

Our analysis stresses the competitive and complementary pressures seen by markets with intermediaries. The zero-profit assumption is ubiquitous when analyzing competitive market organizations and, like here, has been noted to imply cross-cutting implications for efficiency (Mankiw and Whinston, 1986). Whereas Mankiw and Whinston (1986) identify a tendency for over-entry into production markets, we stress under entry. Our framework introduces upstream and downstream complementarities that are typical of many production (or supply-chain) networks. These important complementarities lead to our distinct conclusions. An inefficiency in the market's organization persists, though it is of a different character.

Choi et al. (2014) stress the importance of "critical traders" in network markets. Our analysis complements their conclusion. Equilibrium networks exaggerate the importance of some traders thereby bestowing an abnormal criticality to traders closer to the seller or producer. Likewise, one can interpret the formation of a partnership or other merging behavior as an attempt by traders to bolster their (collective) criticality within the economy as a whole. Such large traders are not only important in an absolute sense, but they also generate indirect market externalities affecting others' profitability. Our model isolates these more subtle channels. Among others, Kranton and Minehart (2000) and Arrow (1975) also examine integration among market participants.

Our model can be extended along many dimensions and incorporated into broader studies of trade with intermediaries. A particularly promising direction concerns developing a more comprehensive understanding of the stability and robustness of networked markets. This is especially salient if traders can form more elaborate network configurations than what we have considered. Similarly, we have focused on a specific market institution, an auction, as mediating exchange. Allowing for alternative or endogenous institutional arrangements—such as consignment contracts, bargaining, or optimal trading mechanisms—among buyers, sellers, and intermediaries, is but one exciting avenue for further analysis.

[29] Wright and Wong (2014) also examine chains of intermediation, though in a search-theoretic context.

A Appendix: Proofs

Lemmas A.1 and A.2 are preliminary results that we use below.

Lemma A.1. *Take an arbitrary trading history and consider trader i in row r.*

1. *Given σ^*_{-i}, trader i cannot earn a positive trading profit in any continuation of the trading game if the asset is held by another trader in row $r-1$ or r.*

2. *Given σ^*_{-i}, the expected resale value of the asset to trader i is $\tilde{\nu}_r = \tilde{\delta}_{r-1}\tilde{b}_{r-1}$ where $\tilde{\delta}_{r-1}$ is the probability assigned by i to the event that there are at least two low-cost agents in row $r-1$ and $\tilde{b}_{r-1}$ is the expected bid of a low-cost trader in row $r-1$.*

Proof. We adopt the convention that buyers are "low-cost agents" in row 0 who bid v. The proof is by induction on r.

Base Case Let $r = 1$. (1) Suppose that the asset is held by an agent in row zero. It is not available for trade and i cannot earn further trading profits. If the asset is sold by another trader in row 1, all agents in row 0 bid v. If i purchases the asset, he must pay at least v. Given σ^*_{-i}, he will be able to resell it only to a buyer at price v. On net, this buy-sell transaction yields zero trading profit. (2) As the trader receives payment only if a buyer in row 0 acquires the asset and further trading profit is not possible, the expected resale value is $\tilde{\nu}_1 = \tilde{\delta}_0 \cdot \tilde{b}_0 = 1 \cdot v = v$.

Induction Hypothesis A trader in row k cannot earn a positive trading profit in any continuation of the trading game if the asset is held by another trader in rows $k-1$ or k. Moreover, the expected resale value of the asset to a trader in row k is $\tilde{\nu}_k = \tilde{\delta}_{k-1}\tilde{b}_{k-1}$.

Inductive Step The base case ($k = 1$) satisfies the induction hypothesis. Therefore, suppose the hypothesis is true for $k = r-1$. We will verify that it is true for $k = r$.

(1) Suppose the asset is being sold by a trader in row $r-1$ or row r. If i is to earn a positive trading profit, he must be able to earn a positive trading profit in at least one of the transactions sketched in Figure A.1.[30]

(A) *Agent i buys the asset from j in row $r-1$ and resells it to k in row $r-1$.* Let b be the payment made by i. It equals the highest bid submitted by an agent in row $r-2$. (If all

[30] Given σ^*_{-i}, the asset will never reach to a row $r'' \geq r+1$. Once the asset reaches row $r'' \leq r-2$, i will not have the opportunity to purchase it.

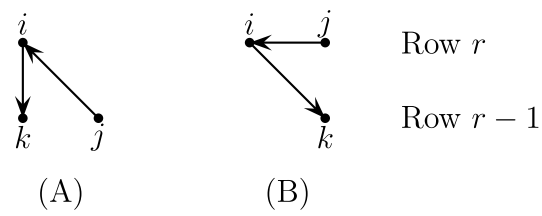

Figure A.1: Potentially profitable transactions.

traders in row $r-2$ bid ℓ, $b=0$.) Suppose i resells the asset. In that auction, bidders in rows $r+1$ and r bid ℓ. By σ^*_{r-1} and the induction hypothesis, low-cost agents in row $r-1$ bid $\tilde{b}_{r-1} = \tilde{\nu}_{r-1} = \tilde{\delta}_{r-2}\tilde{b}_{r-2} \leq b$.[31] Thus, i is unable to resell the asset at a price that yields a strict profit.

(B) Agent i buys the asset from j in row r and resells it to k in row $r-1$. Let b equal the payment made by i. It equals the highest bid submitted by an agent in row $r-1$. (If all traders in row $r-1$ bid ℓ, $b=0$.) When i resells the asset, the maximal submitted bid by row $r-1$ agents is again b. Therefore, the resale price is bounded above by b and i cannot earn a strict profit.

(2) Since bidder i is unable to earn additional trading profit once the asset reaches row $r-1$, his expected resale value is determined by the bid of row $r-1$ low-cost traders. Payment is received only if there are at least two low-cost traders in row $r-1$. Thus, $\tilde{\nu}_r = \tilde{\tilde{\delta}}_{r-1}\tilde{b}_{r-1}$. □

Remark A.1 (Expected Resale Values). Given σ^*_{-i}, we can compute via induction the expected resale value of agent i in row r to be $\tilde{\nu}_r = \prod_{k=1}^{r-1} \tilde{\delta}_k v$.

Remark A.2 (Beliefs). We will argue that the strategy profile outlined in Theorem 1 is supported as an equilibrium by the following belief system. On the equilibrium path, beliefs evolve according to Bayes' rule conditional on the defined strategy profile. In off-equilibrium path situations we specify beliefs as follows:

1. If an agent has not bid in any auction, others maintain their prior beliefs concerning the agent's type.

[31] No new information regarding the resale value of the asset to agents in row $r-2$ was revealed in the interim; therefore, their bids are unchanged.

2. If an agent bids ℓ in the *first* auction in which he participates, in all continuation histories of the trading game others believe this agent has high trading cost.

3. If an agent places a competitive bid (i.e. a bid other than ℓ) in the *first* auction in which he participates, in all continuation histories of the trading game others believe this agent has low trading cost.

Lemma A.2. *Let* $\mu(n) = 1 - (1-p)^n$ *and* $\delta(n) = 1 - (1-p)^n - np(1-p)^{n-1}$.

1. $\mu(0) = 0$ *and* $\mu(n)/p \geq 1$ *for* $n \geq 1$.

2. $\lim_{p\to 0} \frac{\mu(n_k)}{p} = n_k$.

3. *For* $n \geq 1$ *and* $p \in (0,1)$, $\mu(n-1) \geq \delta(n) \geq p\mu(n-1)$.

Proof. (1) $\mu(0) = 1 - (1-p)^0 = 1 - 1 = 0$. Furthermore, since $\mu(n)$ is increasing in n, $\mu(n) = 1 - (1-p)^n \geq 1 - (1-p)^1 = p \implies \mu(n)/p \geq 1$. (2) Applying l'Hôpital's Rule, $\lim_{p\to 0} \frac{\mu(n)}{p} = \lim_{p\to 0} \frac{1-(1-p)^n}{p} = \lim_{p\to 0} \frac{n(1-p)^{n-1}}{1} = n$. (3) A direct calculation shows that $\mu(n-1) - \delta(n) = (n-1)p(1-p)^{n-1} \geq 0$ for $n \geq 1$. Similarly,

$$\delta(n) - p\mu(n-1) = 1 - p - (1-p)^{n-1}(1 + (n-2)p).$$

Thus, $\delta(1) - p\mu(0) = 0$ and $\delta(2) - p\mu(1) = 0$. On the other hand, if $n \geq 2$ we see that

$$\frac{d}{dn}\left[\delta(n) - p\mu(n-1)\right] = -(1-p)^{n-1}(p + (1 + (n-2)p)\log(1-p)) \geq 0.$$

Thus, $\delta(n) - p\mu(n-1) \geq 0$ as required. ◻

Proof of Theorem 1. Consider trader i in row r. There are two cases depending on the asset's trading history. First, suppose the asset is sold by a trader in row $r-1$ or row r. By Lemma A.1 the expected additional trading profit of agent i is zero. Therefore, the bid ℓ is optimal. Suppose instead that the asset is sold by a trader in row $r+1$. If agent i successfully acquires the asset, given σ^*_{-i} the asset's expected resale value is $\tilde{\nu}_r \leq v$. Thus, if i has high trading cost, ℓ is the optimal bid. If i has low trading costs, an argument parallel to that confirming that "bidding one's valuation" is optimal in a second-price auction (Vickrey, 1961) confirms that $\tilde{\nu}_r$ is an optimal bid. ◻

Remark A.3. If $\mathbf{n}$ is the network's configuration, equilibrium-path expected resale values are $\nu_r = \Pi_{k=1}^{r-1}\delta(n_k) \cdot v$. Equilibrium-path bids of low-cost traders are $b_r = \nu_r$.

Corollary A.1. *Consider the equilibrium defined by Theorem 1. The ex ante expected profits of a trader in row r is* $\pi_r(\mathbf{n}) = \prod_{k=1}^{r-1} \delta(n_k) \times \prod_{k=r+1}^{R} \mu(n_k) \times p \times (1-p)^{n_r-1} v$.

Proof. For the asset to reach row $r+1$, at least one trader in each row $k \geq r+1$ must have low trading costs. This event occurs with probability $\prod_{k=r+1}^{R} \mu(n_k)$. With probability p agent i in row r will have low trading cost and will bid ν_r in equilibrium. With probability $(1-p)^{n_r-1}$ all other traders in row r have a high trading cost and i acquires the asset for a price of zero. With probability $1-(1-p)^{n_r-1}$, at least one other trader in row r also has a low trading cost and similarly bids ν_r. Hence, i either does not acquire the asset or must pay ν_r. Thus, the expected surplus to i is $(1-p)^{n_r-1}(\nu_r - 0)$. Since $\nu_r = \prod_{k=1}^{r-1} \delta(n_k) \cdot v$, combining the preceding observations yields the conclusion. □

Proof of Theorem 2. Without loss of generality we can assume $c_v = 0$. First, consider a partnership $\mathbf{m}$ where $m_r \geq 2$ for some r. Since $\lim_{p\to 0} \delta(n) = 0$ and $\lim_{p\to 0} \mu(n) = 0$, $\lim_{p\to 0} \pi_{\mathbf{m}} = 0$. Therefore, there exists a p sufficiently small such that $\pi_{\mathbf{m}}(\mathbf{n}) - \zeta(\mathbf{m}) \leq \pi_{\mathbf{m}}(\mathbf{n}) - c_h < 0 \leq \sum_r m_r \pi_r(\mathbf{n})$. Hence, the network is stable.

Henceforth, we need only consider partnerships where $m_r \leq 1$ for all r. Thus, $\zeta(\mathbf{m}) = 0$ since $c_v = 0$. Furthermore, we can assume that $\underline{m} < \bar{m}$. There are several cases depending on the underlying network structure.

1. If $n_r = 1$ for some $r \leq \underline{m} - 1$, then $\prod_{k=1}^{\underline{m}-1} \delta(n_k) = 0$. Therefore, $\pi_{\mathbf{m}}(\mathbf{n}) = 0$. Thus, a partnership is not profitable.

2. Suppose $n_r \geq 2$ for all $r \leq \underline{m} - 1$ and $n_{\underline{m}} = 1$. In this case,

$$\pi_{\underline{m}}(\mathbf{n}) = \prod_{k=1}^{\underline{m}-1} \delta(n_k) \cdot \prod_{k=\underline{m}+1}^{R} \mu(n_k) \cdot p$$

and

$$\pi_{\mathbf{m}}(\mathbf{n}) = \prod_{k=1}^{\underline{m}-1} \delta(n_k) \cdot \prod_{k=\bar{m}+1}^{R} \mu(n_k) \cdot \prod_{k=\underline{m}}^{\bar{m}} \mu(m_k).$$

Therefore, since $\mu(m_{\underline{m}}) = p$ and $\mu(n_k) \geq \mu(m_k)$,

$$\frac{\pi_{\underline{m}}(\mathbf{n})}{\pi_{\mathbf{m}}(\mathbf{n})} = \frac{p \prod_{k=\underline{m}+1}^{\bar{m}} \mu(n_k)}{p \prod_{k=\underline{m}+1}^{\bar{m}} \mu(m_k)} \geq 1.$$

Thus, $\pi_{\mathbf{m}}(\mathbf{n}) \leq \pi_{\underline{m}}(\mathbf{n}) \leq \sum_r m_r \pi_r(\mathbf{n})$. Thus, the partnership is not sufficiently profitable.

3. Suppose $n_k \geq 2$ for all $k \leq \underline{m}$ but $n_{\underline{m}+1} = 1$. (This implies $n_{\underline{m}+1} = m_{\underline{m}+1} = 1$.) In this case,

$$\pi_{\underline{m}}(\mathbf{n}) = \prod_{k=1}^{\underline{m}-1} \delta(n_k) \cdot \prod_{k=\underline{m}+2}^{R} \mu(n_k) \cdot \left[\mu(n_{\underline{m}+1}) \cdot p \cdot (1-p)^{n_{\underline{m}}-1}\right]$$

and

$$\pi_{\underline{m}+1}(\mathbf{n}) = \prod_{k=1}^{\underline{m}-1} \delta(n_k) \cdot \prod_{k=\underline{m}+2}^{R} \mu(n_k) \cdot \left[\delta(n_{\underline{m}}) \cdot p \cdot (1-p)^{n_{\underline{m}+1}-1}\right].$$

Since $n_{\underline{m}+1} = m_{\underline{m}+1} = 1$,

$$\pi_{\mathbf{m}}(\mathbf{n}) = \prod_{k=1}^{\underline{m}-1} \delta(n_k) \cdot \prod_{k=\bar{m}+1}^{R} \mu(n_k) \cdot \prod_{k=\underline{m}+2}^{\bar{m}} \mu(m_k) \cdot \prod_{k=\underline{m}}^{\underline{m}+1} \mu(m_k).$$

Therefore,

$$\begin{aligned}
\frac{\pi_{\underline{m}}(\mathbf{n}) + \pi_{\underline{m}+1}(\mathbf{n})}{\pi_{\mathbf{m}}(\mathbf{n})} &= \left[\prod_{k=\underline{m}+2}^{\bar{m}} \frac{\mu(n_k)}{p}\right] \cdot \left[\frac{p^2(1-p)^{n_{\underline{m}}-1} + \delta(n_{\underline{m}})p}{p \cdot p}\right] \\
&= \left[\prod_{k=\underline{m}+2}^{\bar{m}} \frac{\mu(n_k)}{p}\right] \cdot \left[(1-p)^{n_{\underline{m}}-1} + \frac{\delta(n_{\underline{m}})}{p}\right]
\end{aligned}$$

From Lemma A.2, $\frac{\mu(n_k)}{p} \geq 1$. Moreover, also from Lemma A.2

$$\begin{aligned}
\delta(n_{\underline{m}}) \geq p\mu(n_{\underline{m}} - 1) &\implies \delta(n_{\underline{m}}) \geq p - p(1-p)^{n_{\underline{m}}-1} \\
&\implies (1-p)^{n_{\underline{m}}-1} + \frac{\delta(n_{\underline{m}})}{p} \geq 1.
\end{aligned}$$

Hence, $\frac{\pi_{\underline{m}}(\mathbf{n})+\pi_{\underline{m}+1}(\mathbf{n})}{\pi_{\mathbf{m}}(\mathbf{n})} \geq 1$ and thus $\pi_{\mathbf{m}}(\mathbf{n}) \leq \pi_{\underline{m}}(\mathbf{n}) + \pi_{\underline{m}+1}(\mathbf{n}) \leq \sum_r m_r \pi_r(\mathbf{n})$.

4. Suppose $n_k \geq 2$ for all $k \leq \underline{m} + 1$. In this case,

$$\pi_{\underline{m}}(\mathbf{n}) = \prod_{k=1}^{\underline{m}-1} \delta(n_k) \cdot \prod_{k=\underline{m}+1}^{R} \mu(n_k) \cdot p \cdot (1-p)^{n_{\underline{m}}-1}$$

and

$$\pi_{\mathbf{m}}(\mathbf{n}) = \prod_{k=1}^{\underline{m}-1} \delta(n_k) \cdot \prod_{k=\bar{m}+1}^{R} \mu(n_k) \cdot p^{\bar{m}-\underline{m}+1} \cdot \left(1 - \mu(n_{\bar{m}} - 1) \prod_{k=\underline{m}}^{\bar{m}-1} \delta(n_k - 1)\right).$$

Thus,

$$\frac{\pi_{\underline{m}}(\mathbf{n})}{\pi_{\mathbf{m}}(\mathbf{n})} = \left[\prod_{k=\underline{m}+1}^{\bar{m}} \frac{\mu(n_k)}{p}\right] \cdot \frac{p}{p} \cdot \underbrace{\left[\frac{(1-p)^{n_{\underline{m}}-1}}{1 - \mu(n_{\bar{m}} - 1) \prod_{k=\underline{m}}^{\bar{m}-1} \delta(n_k - 1)}\right]}_{[1]}.$$

As $p \to 0$, term [1] converges to 1 and by Lemma A.2, $\lim_{p \to 0} \prod_{k=\underline{m}+1}^{\bar{m}} \frac{\mu(n_k)}{p} = \prod_{k=\underline{m}+1}^{\bar{m}} n_k \geq 2$. Therefore, $\lim_{p \to 0} \frac{\pi_{\underline{m}}(\mathbf{n})}{\pi_{\mathbf{m}}(\mathbf{n})} > 1$ and for p sufficiently small, $\pi_{\mathbf{m}}(\mathbf{n}) \leq \pi_{\underline{m}}(\mathbf{n}) \leq \sum_r m_r \pi_r(\mathbf{n})$.

For every feasible partnership, the above argument has confirmed that there exists a $p > 0$ sufficiently small such that $\sum_r m_r \pi_r(\mathbf{n}) \geq \pi_{\mathbf{n}}(\mathbf{n}) - \zeta(\mathbf{m})$. Since there is a finite number of possible partnerships, there exists a $\hat{p} > 0$ sufficiently small such that the underlying network $\mathbf{n}$ is stable. □

Proof of Theorem 3. In a network without a partnership, the expected profit of a row-1 trader is

$$\pi_1(\mathbf{n}) = \prod_{k=2}^{\bar{m}} \mu(n_k) \cdot p \cdot (1-p)^{n_1 - 1}.$$

If two row-1 traders merge, i.e. $\mathbf{m} = (2, 0, \ldots)$, the partnership's expected profit is

$$\pi_{\mathbf{m}}(\mathbf{n}) = \prod_{k=2}^{\bar{m}} \mu(n_k) \cdot (1 - (1-p)^2) \cdot (1-p)^{n_1 - 2}$$

Then,

$$\begin{aligned}
\pi_{\mathbf{m}}(\mathbf{n}) > 2\pi_1(\mathbf{n}) &\iff (1 - (1-p)^2)(1-p)^{n_1-2} > 2p(1-p)^{n_1-1} \\
&\iff (1-p)^{2+n_1} p^2 > 0,
\end{aligned}$$

which holds for all $p \in (0, 1)$. Hence, the proposed merger is profitable when $c_h = 0$. □

Proof of Theorem 4. Let $\mathbf{n}$ be a network configuration such that $n_r \geq 1$ for all r. Noting that $n_r p(1-p)^{n_r - 1} = \mu(n_r) - \delta(n_r)$ and that $\mu(n_r) \neq 0$, we can compute the sum of intermediary traders' expected profits to be

$$\begin{aligned}
\sum_{r=1}^{R} n_r \pi_r(\mathbf{n}) &= \sum_{r=1}^{R} n_r \left[\prod_{k=1}^{r-1} \delta(n_k)\right] \left[p\,(1-p)^{n_r-1}\right] \left[\prod_{k=r+1}^{R} \mu(n_k)\right] \\
&= \sum_{r=1}^{R} \left[\prod_{k=1}^{r-1} \delta(n_k)\right] \left[\mu\,(n_r) - \delta\,(n_r)\right] \left[\prod_{k=r+1}^{R} \mu(n_k)\right] \\
&= \left[\prod_{k=1}^{R} \mu(n_k)\right] \sum_{r=1}^{R} \left(\prod_{k=1}^{r-1} \frac{\delta\,(n_k)}{\mu(n_k)} - \prod_{k=1}^{r} \frac{\delta(n_k)}{\mu(n_k)}\right) \\
&= \left[\prod_{k=1}^{R} \mu(n_k)\right] \left(1 - \prod_{k=1}^{R} \frac{\delta(n_k)}{\mu\,(n_k)}\right) \\
&= \prod_{k=1}^{R} \mu(n_k) - \prod_{k=1}^{R} \delta(n_k)
\end{aligned}$$

The expected profits of the seller are $\pi_{R+1}(\mathbf{n}) = \prod_{k=1}^{R} \delta(n_k)$. Buyers' expected welfare is zero. Hence, $\sum_{r=1}^{R} n_r \pi_r(\mathbf{n}) + \pi_{R+1}(\mathbf{n}) = \prod_{k=1}^{R} \mu(n_k)$. □

Proof of Theorem 5. (1) From (2), $\pi_r(\mathbf{n}^*) \le p(1-p)^{n_r-1}$; therefore,

$$\pi_r(\mathbf{n}^*) - \kappa \ge 0 \implies p(1-p)^{n_r^*-1} - \kappa \ge 0 \implies n_r^* \le \bar{n} = \left\lceil 1 + \frac{\log(\kappa) - \log(p)}{\log(1-p)} \right\rceil .$$

(2) Necessity follows from the definition of equilibrium and part (1). To show sufficiency, we define a tâtonnement-style mapping that converges to an equilibrium. First, choose $\mathbf{n}^0$ such that $\pi_r(\mathbf{n}^0) - \kappa \ge 0$ for all r. Define $\mathscr{Q}_r(\mathbf{n}) = \{\tilde{n}_r \in \mathbb{N} \colon \pi_r(\tilde{n}_r, \mathbf{n}_{-r}) - \kappa \ge 0, n_r \le \tilde{n}_r \le \bar{n}\}$ and let $\hat{n}_r = \max \mathscr{Q}_r(\mathbf{n})$. Next, define $A_r(\cdot)$ as

$$A_r(\mathbf{n}) = \begin{cases} (\hat{n}_r, \mathbf{n}_{-r}) & \text{if } \mathscr{Q}_r(\mathbf{n}) \neq \emptyset \\ \mathbf{n}^0 & \text{if } \mathscr{Q}_r(\mathbf{n}) = \emptyset \end{cases}$$

Thus, given $\mathbf{n}$, $A_r(\cdot)$ increases n_r until adding another agent to row r (holding $\mathbf{n}_{-r}$ fixed) yields negative profits. Composing these mappings together gives

$$A(\mathbf{n}) = (A_1 \circ \cdots \circ A_R)(\mathbf{n}). \tag{A.1}$$

We argue that A has a fixed point, $A(\mathbf{n}^*) = \mathbf{n}^*$, and that $\mathbf{n}^*$ is an equilibrium.

To show that A has a fixed point we first establish that if $\pi_r(\mathbf{n}) - \kappa \geq 0$ for all r, then $A(\mathbf{n}) \geq \mathbf{n}$. Suppose $\pi_R(\mathbf{n}) - \kappa \geq 0$. Then $\mathcal{Q}_R(\mathbf{n}) \neq \emptyset$. So, $A_R(\mathbf{n}) \geq \mathbf{n}$ since n_R may have increased. Now consider any r and let $\tilde{\mathbf{n}} = (n_1, \ldots, n_r, \tilde{n}_{r+1}, \ldots, \tilde{n}_R)$ where the first r terms are unchanged relative to $\mathbf{n}$ and $(\tilde{n}_{r+1}, \ldots, \tilde{n}_R) \geq (n_{r+1}, \ldots, n_R)$. Then $\pi_r(n_r, \tilde{\mathbf{n}}_{-r}) - \kappa \geq \pi_r(n_r, \mathbf{n}_{-r}) - \kappa \geq 0$. Therefore, $A_r(\tilde{\mathbf{n}}) \geq \tilde{\mathbf{n}}$. This implies $A(\mathbf{n}) \geq \mathbf{n}$. Note also that for all r, $\pi_r(A(\mathbf{n})) - \kappa \geq 0$. Indeed, if we let $\tilde{\mathbf{n}} = A(\mathbf{n})$, we see that

$$\pi_r(\tilde{\mathbf{n}}) - \kappa \geq \pi_r(n_1, \ldots, n_{r-1}, \tilde{n}_r, \ldots, \tilde{n}_R) - \kappa \geq 0.$$

Finally, consider the sequence $\mathbf{n}^{t+1} = A(\mathbf{n}^t)$ starting at $\mathbf{n}^0$. $\mathbf{n}^t$ is a non-decreasing sequence and for each t, $\pi_r(\mathbf{n}^t) - \kappa \geq 0$. Since $\mathbf{n}^t$ is bounded by $(\bar{n}, \ldots, \bar{n})$, the sequence $\{\mathbf{n}^t\}$ converges to a limit $\mathbf{n}^*$. Thus, there exists a configuration such that $\mathbf{n}^* = A(\mathbf{n}^*)$.

Take $\mathbf{n}^* = A(\mathbf{n}^*)$ and suppose that $\mathbf{n}^*$ is not an equilibrium. Therefore, there exists some row $\hat{r}$ such that either (1) $\pi_{\hat{r}}(\mathbf{n}^*) - \kappa < 0$ or (2) $\pi_{\hat{r}}(n^*_{\hat{r}} + 1, \mathbf{n}^*_{-\hat{r}}) - \kappa \geq 0$. We address both cases.

1. Suppose that $\pi_{\hat{r}}(\mathbf{n}^*) - \kappa < 0$. Then, $A_{\hat{r}}(\mathbf{n}^*) = \mathbf{n}^0$ since $\mathcal{Q}_{\hat{r}}(\mathbf{n}^*) = \emptyset$. Therefore $\mathbf{n}^* = (n^*_1, \ldots, n^*_{\hat{r}-1}, n^0_{\hat{r}}, \ldots, n^0_R)$. Thus, recalling that $\pi_r(n_r, \mathbf{n}_{-r})$ is increasing in $\mathbf{n}_{-r}$ and $\mathbf{n}^* \geq \mathbf{n}^0$,

$$\pi_{\hat{r}}(\mathbf{n}^*) - \kappa = \pi_{\hat{r}}(n^*_1, \ldots, n^*_{\hat{r}-1}, n^0_{\hat{r}}, \ldots, n^0_R) - \kappa \geq \pi_{\hat{r}}(\mathbf{n}^0) - \kappa \geq 0,$$

which is a contradiction.

2. Suppose instead that $\pi_{\hat{r}}(n^*_{\hat{r}} + 1, \mathbf{n}^*_{-\hat{r}}) - \kappa \geq 0$. But then, from the definition of $\mathcal{Q}_{\hat{r}}$, $n^*_{\hat{r}} + 1 \in \mathcal{Q}_r(\mathbf{n}^*)$. This implies $n^*_{\hat{r}} \geq n^*_{\hat{r}} + 1$, which is a contradiction.

Therefore $\mathbf{n}^* = A(\mathbf{n}^*)$ is an equilibrium configuration. □

Proof of Theorem 6. Let $\mathbf{n}$ and $\mathbf{n}'$ be equilibria. Choose r and without loss of generality suppose $n_r \geq n'_r$. Then, $\pi_r(\mathbf{n} \vee \mathbf{n}') - \kappa = \pi_r(n_r, \mathbf{n}_{-r} \vee \mathbf{n}'_{-r}) - \kappa \geq \pi_r(n_r, \mathbf{n}_{-r}) - \kappa \geq 0$. Applying the mapping $A(\cdot)$ as in the proof of Theorem 5 but with $\mathbf{n} \vee \mathbf{n}'$ as the initial condition allows us to construct a sequence of configurations converging to an equilibrium, say $\mathbf{x}$. Since the sequence is increasing, $\mathbf{x} \geq \mathbf{n} \vee \mathbf{n}'$. □

Proof of Corollary 1. By Theorem 6, the set of equilibria are a directed set. This set is finite. The conclusion follows. □

Proof of Theorem 7. We argue by contradiction. Suppose $\mathbf{n}^*$ is an equilibrium such that for some $1 \leq r \leq R-1$, $n_r^* < n_{r+1}^*$. Since $\mathbf{n}^*$ is an equilibrium, the following inequalities hold:

$$\prod_{k=1}^{r-1} \delta(n_k^*) \left[p(1-p)^{n_r^*-1}\mu(n_{r+1}^*)\right] \prod_{k=r+2}^{R} \mu(n_k^*) \geq \kappa > \prod_{k=1}^{r-1} \delta(n_k^*) \left[p(1-p)^{n_r^*}\mu(n_{r+1}^*)\right] \prod_{k=r+2}^{R} \mu(n_k^*)$$
$$\prod_{k=1}^{r-1} \delta(n_k^*) \left[\delta(n_r^*)p(1-p)^{n_{r+1}^*-1}\right] \prod_{k=r+2}^{R} \mu(n_k^*) \geq \kappa > \prod_{k=1}^{r-1} \delta(n_k^*) \left[\delta(n_r^*)p(1-p)^{n_{r+1}^*}\right] \prod_{k=r+2}^{R} \mu(n_k^*)$$

To simplify, let $\tilde{\kappa} \equiv \kappa / (\prod_{k=1}^{r-1} \delta(n_k^*) \times \prod_{k=r+2}^{R} \mu(n_k^*))$, then the above inequalities become

$$p(1-p)^{n_r^*-1}\mu(n_{r+1}^*) \geq \tilde{\kappa} > p(1-p)^{n_r^*}\mu(n_{r+1}^*)$$
$$\delta(n_r^*)p(1-p)^{n_{r+1}^*-1} \geq \tilde{\kappa} > \delta(n_r^*)p(1-p)^{n_{r+1}^*}$$

From these inequalities, we see that $\delta(n_r^*)(1-p)^{n_{r+1}^*-1} > (1-p)^{n_r^*}\mu(n_{r+1}^*)$. However, since $n_{r+1}^* \geq n_r^* + 1$, $(1-p)^{n_{r+1}^*-1} \leq (1-p)^{n_r^*}$. Similarly, $\delta(n_r^*) \leq \delta(n_r^*+1) \leq \delta(n_{r+1}^*) < \mu(n_{r+1}^*)$. As the preceding terms are all non-negative, $\delta(n_r^*)(1-p)^{n_{r+1}^*-1} < (1-p)^{n_r^*}\mu(n_{r+1}^*)$, which is a contradiction. □

Proof of Theorem 8. Suppose that a solution to (OPT) is such that $\hat{n}_r > \hat{n}_{r'} \geq 1$ for some r and r'. Hence,

$$\prod_{k=1}^{R} \mu(\hat{n}_k) - \kappa \sum_{k=1}^{R} \hat{n}_k \geq \prod_{k \neq r} \mu(\hat{n}_k) \cdot \mu(\hat{n}_r - 1) - \kappa \sum_{k \neq r} \hat{n}_k - \kappa(\hat{n}_r - 1). \quad \text{(A.2)}$$

Since $\mu(n)$ is concave and nondecreasing, $\mu(\hat{n}_r) - \mu(\hat{n}_r - 1) \leq \mu(\hat{n}_{r'} + 1) - \mu(\hat{n}_{r'})$. Thus, rearranging terms in (A.2) and substituting gives

$$
\begin{aligned}
\text{(A.2)} \implies & \prod_{k\neq r,r'} \mu(\hat{n}_k) \cdot \mu(\hat{n}_{r'}) \big[\mu(\hat{n}_r) - \mu(\hat{n}_r - 1)\big] \geq \kappa \\
\implies & \prod_{k\neq r,r'} \mu(\hat{n}_k) \cdot \mu(\hat{n}_{r'}) \big[\mu(\hat{n}_{r'} + 1) - \mu(\hat{n}_{r'})\big] \geq \kappa \\
\implies & \prod_{k\neq r,r'} \mu(\hat{n}_k) \cdot \mu(\hat{n}_{r}) \big[\mu(\hat{n}_{r'} + 1) - \mu(\hat{n}_{r'})\big] > \kappa \\
\implies & \prod_{k\neq r'} \mu(\hat{n}_k) \cdot \mu(\hat{n}_{r'} + 1) - \kappa > \prod_{k=1}^{R} \mu(\hat{n}_k) \\
\implies & \prod_{k\neq r'} \mu(\hat{n}_k) \cdot \mu(\hat{n}_{r'} + 1) - \kappa - \kappa \sum_{k=1}^{R} \hat{n}_k > \prod_{k=1}^{R} \mu(\hat{n}_k) - \kappa \sum_{k=1}^{R} \hat{n}_k.
\end{aligned}
$$

The final expression contradicts $\hat{\mathbf{n}}$ being a solution to (OPT).

To show the theorem's second part, let $\hat{\mathbf{n}} = (\hat{n}, \ldots, \hat{n})$ solve (OPT) and let $\mathbf{n}^*$ be an equilibrium configuration. To work toward a contradiction, suppose $n_1^* > \hat{n}$. Let $\bar{r} = \max\{r : n_r^* = n_1^*\}$. Since $\mathbf{n}^*$ is an equilibrium,

$$
\begin{aligned}
\pi_{\bar{r}}(\mathbf{n}^*) \geq \kappa \implies & \prod_{k>\bar{r}} \mu(n_k^*) \cdot \prod_{k<\bar{r}} \delta(n_k^*) \cdot p(1-p)^{n_{\bar{r}}^*-1} \geq \kappa \\
\implies & \prod_{k>\bar{r}} \mu(n_1^* - 1) \cdot \prod_{k<\bar{r}} \mu(n_1^* - 1) \cdot p(1-p)^{n_1^*-1} \geq \kappa \\
\implies & \mu(n_1^* - 1)^{R-1} \left(\mu(n_1^*) - \mu(n_1^* - 1)\right) \geq \kappa \qquad \text{(A.3)}
\end{aligned}
$$

The first implication is from the definition of $\pi_r(\mathbf{n}^*)$. The second implication follows since $n_1^* - 1 \geq n_k^*$ for all $k > \bar{r}$, $n_k^* = n_1^*$ for all $k \leq \bar{r}$, and $\mu(n-1) \geq \delta(n)$ for all $n \geq 1$ by Lemma A.2. The final implication follows from a regrouping of terms and the substitution $p(1-p)^{n_1^*-1} = \mu(n_1^*) - \mu(n_1^* - 1)$.

Consider the following difference written as a telescoping sum:

$$
\mu(n_1^*)^R - \mu(n_1^* - 1)^R = \sum_{k=0}^{R-1} \left[\mu(n_1^* - 1)^{R-1-k} \mu(n_1^*)^{k+1} - \mu(n_1^* - 1)^{R-k} \mu(n_1^*)^{k}\right].
$$

Examining each term in the sum shows

$$\begin{aligned}\mu(n_1^*-1)^{R-1-k}\mu(n_1^*)^{k+1}-\mu(n_1^*-1)^{R-k}\mu(n_1^*)^{k} &= \mu(n_1^*-1)^{R-1-k}\mu(n_1^*)^{k}(\mu(n_1^*)-\mu(n_1^*-1)) \\ &\geq \mu(n_1^*-1)^{R-1}\left(\mu(n_1^*)-\mu(n_1^*-1)\right) \\ &\geq \kappa.\end{aligned}$$

The final inequality follows from (A.3). Hence,

$$\mu(n_1^*)^R-\mu(n_1^*-1)^R \geq R\kappa. \tag{A.4}$$

Recall the welfare-maximizing configuration $\hat{\mathbf{n}}$. Since $\hat{n}_r=\hat{n}$ for all r, $\hat{n}$ must also solve $\max_{n\in\mathbb{Z}_+}\mu(n)^R-Rn\kappa$. This objective function is single-peaked and its greatest solution must satisfy the following "discretized first-order condition": $\mu(\hat{n})^R-\mu(\hat{n}-1)^R\geq R\kappa>\mu(\hat{n}+1)^R-\mu(\hat{n})^R$. Since $n\mapsto\mu(n+1)^R-\mu(n)^R$ is a decreasing function and $n_1^*>\hat{n}$, the preceding inequality implies that $R\kappa>\mu(\hat{n}+1)^R-\mu(\hat{n})^R\geq\mu(n_1^*)^R-\mu(n_1^*-1)^R$. But these inequalities contradict (A.4). Thus, $n_1^*\leq\hat{n}$. Since $n_r^*\leq n_1^*$ for all r, $\mathbf{n}^*\leq\hat{\mathbf{n}}$. $\square$

The following appendix is intended for online publication only.

B Exchange in the Presence of a Partnership

Theorem 1 in the main text characterizes exchange in a fixed network. In this supplement we extend Theorem 1 to accommodate a partnership.

Let $\mathbf{n} = (n_1, \ldots, n_R)$ represent a trading network where there are $n_r \geq 1$ agents in row r. Let $\mathbf{m} = (m_1, \ldots, m_R)$ be a partnership with m_r members in row r. Agents not belonging to the partnership are *independent traders*. Let $\underline{m} = \min\{r : m_r \geq 1\}$ and $\bar{m} = \max\{r : m_r \geq 1\}$. For example, consider Figure B.1, which reproduces Figure 3 from the main text. Network A is the trading network $\mathbf{n} = (4, 4, 3, 2)$. Network B modifies Network A by introducing the partnership $\mathbf{m} = (0, 2, 1, 0)$. Thus, $\underline{m} = 2$ and $\bar{m} = 3$.

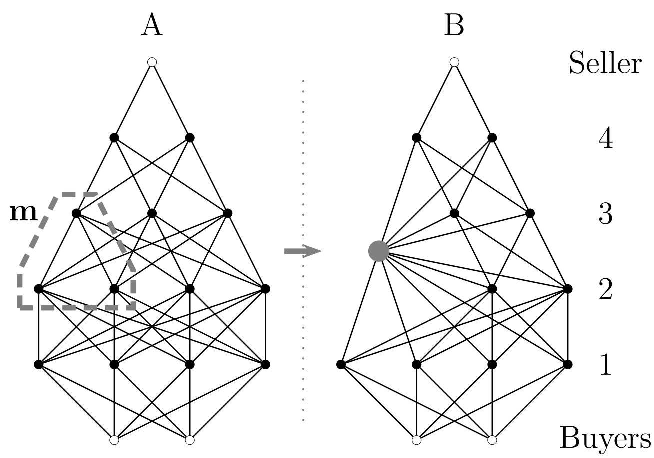

Figure B.1: The formation of the partnership $\mathbf{m} = (0, 2, 1, 0)$. (Within-row links are not illustrated for clarity.)

Remark B.1. If $\bar{m} = \underline{m}$, then Theorem 1 applies with minimal modifications. The single-row partnership is like a single trader in row $\bar{m}$. It has a low trading cost with probability $p_{\mathbf{m}}$. We henceforth assume that $\underline{m} < \bar{m}$.

B.1 Notation and Terminology

We rely on some specific notation.

- σ^* — the strategy profile defined in Theorem B.1 below. σ^*_r denotes the strategy of an independent trader in row r while $\sigma^*_{\mathbf{m}}$ denotes the partnership's strategy. σ^*_i denotes the strategy of a particular independent trader i. (Due to context, confusion between σ^*_r, σ^*_i, and $\sigma^*_{\mathbf{m}}$ should not result.) $\sigma^*_{-\mathbf{m}}$ and σ^*_{-i} have their usual meanings as a strategy profile removing $\sigma^*_{\mathbf{m}}$ or σ^*_i, respectively.

- $\tilde{b}_r$ — the expected bid of a low-cost, independent trader in row r given σ^*_r.

- $\tilde{\nu}_r$ ($\tilde{\nu}_{\mathbf{m}}$) — the expected resale value of the asset to an independent trader in row r (the partnership).

- $\tilde{\mu}_r$ — the probability assigned by agents in row $r' \neq r$ (or by the partnership) to the event that there is *at least one* low-cost, independent trader in row r.

  When an agent holds his prior belief concerning this value, $\tilde{\mu}_r = \mu(n_r - m_r)$ where $\mu(n) = 1 - (1-p)^n$.

- $\tilde{\delta}_r$ — the probability assigned by agents in row $r' \neq r$ (or by the partnership) to the event that there are *at least two* low-cost independent traders in row r.

  When an agent holds his prior belief concerning this value, $\tilde{\delta}_r = \delta(n_r - m_r)$ where $\delta(n) = 1 - (1-p)^n - np(1-p)^{n-1}$.

Unless noted otherwise, expectations are conditional on the strategy adopted by other agents, on the trading history, and given the specification of on- and off-equilibrium path beliefs (see Section B.3). We call bids $b \neq \ell$ competitive bids.

B.2 Main Theorem

The following theorem describes an equilibrium of the trading game in the presence of a partnership. Beliefs supporting the defined strategy profile as a perfect Bayesian equilibrium are specified in Section B.3. This theorem generalizes Theorem 1 from the main text.

Theorem B.1. *Fix a trading network with configuration* $\mathbf{n}$. *Let* $\mathbf{m}$ *be a partnership such that* $\underline{m} < \bar{m}$. *There exists a perfect Bayesian equilibrium of the trading game such that:*

1. *Independent trader* i *in row* $r \leq \underline{m} - 1$ *adopts the following strategy:*

   (a) *If trading costs are low and the asset is being sold by an independent trader in row* $r+1$, *place a bid equal to the asset's expected resale value to agent* i *conditional on all available information and on* σ^*_{-i}. *(Buyers in row 0 bid* v.*)*

(b) *Otherwise, bid ℓ.*

2. *The partnership adopts the following strategy:*

   (a) *If trading costs are low and the asset is being sold by an independent trader in row $r \in \{\underline{m}+1, \ldots, \bar{m}+1\}$, place a bid equal to the asset's expected resale value to the partnership conditional on all available information and on $\sigma^*_{-\mathbf{m}}$.*

   (b) *If the asset is being sold by an agent in rows $\underline{m}$ or $\underline{m}-1$ or trading costs are high, bid ℓ.*

3. *Independent trader i in row $\underline{m} \leq r \leq \bar{m}-1$ adopts the following strategy:*

   (a) *If trading costs are low and the asset is being sold by an independent trader in row $r+1$, place a bid equal to the asset's expected resale value to agent i conditional on all available information and on σ^*_{-i}.*

   (b) *Otherwise, bid ℓ.*

4. *Independent trader i in row $r = \bar{m}$ adopts the following strategy:*

   (a) *If trading costs are low and the asset is being sold by an independent trader in row $r+1$:*

      i. *If the partnership has not yet bid for the asset, place a bid equal to the asset's expected resale value conditional on all available information, σ^*_{-i}, and conditional on the partnership having high trading cost with probability 1.*

      ii. *If the partnership has bid for the asset, place a bid equal to the asset's expected resale value conditional on all available information and σ^*_{-i}.*

   (b) *Otherwise, bid ℓ.*

5. *Independent trader i in row $r \geq \bar{m}+1$ adopts the following*

   (a) *If trading costs are low and the asset is being sold by an independent trader in row $r+1$, place a bid equal to the asset's expected resale value to agent i conditional on all available information and on σ^*_{-i}.*

   (b) *Otherwise, bid ℓ.*

Like Theorem 1, Theorem B.1 specifies via induction a bid for every agent for every trading history. Of course, the expected resale value of the asset depends on agents' beliefs conceding the trading costs of other agents in the market. Below we specify how beliefs evolve.

B.3 Beliefs

We argue that the strategy profile outlined above is supported as an equilibrium by the following beliefs. As usual, on the equilibrium path beliefs evolve according to Bayes' rule conditional on σ^*. In off-equilibrium path situations beliefs are defined as follows:

1. If an agent (an independent trader or the partnership) has not bid in any auction, others maintain their prior beliefs concerning that agent's type.

2. If an independent trader in row $r \notin \{\underline{m}, \dots, \bar{m}-1\}$ bids ℓ in the first auction in which he bids, then in all continuations of the trading history others believe this trader has a high trading cost. If instead this trader places a competitive bid in that auction, then others believe this agent has a low trading cost.

3. If an independent trader in row $r \in \{\underline{m}, \dots, \bar{m}-1\}$ bids ℓ in the first auction in which he bids and that auction is not organized by the partnership, then in all continuations of the trading history others believe this trader has a high trading cost. If instead this trader places a competitive bid in that auction, then others believe this agent has a low trading cost.

4. If an independent trader in row $r \in \{\underline{m}, \dots, \bar{m}-1\}$ ever places a competitive bid in an auction organized by the partnership, then in all continuations of the trading history others believe this trader has a low trading cost. Otherwise, if the trader has only bid ℓ in such auctions, others do not update their beliefs concerning this agent's type.

5. If the partnership bids ℓ in the first auction in which it bids, then in all continuations of the trading history others believe the partnership has a high cost. If the partnership places a competitive bid in that auction, then others believe the partnership has a low trading costs.

B.4 Preliminary Remarks and Lemmas

Remark B.2. Lemma A1 from the main text applies essentially verbatim to all independent traders in rows $r \leq \underline{m}$ and $r \geq \bar{m}+1$. It also applies to the partnership when the asset is held by an independent trader in row $\underline{m}-1$ or $\underline{m}$.

Lemma B.1. *Take an arbitrary trading history and suppose all independent traders follow* $\sigma^*_{-\mathbf{m}}$. *The asset's expected resale value to the partnership is* $\tilde{\nu}_{\mathbf{m}} = \tilde{\delta}_{\underline{m}-1}\tilde{b}_{\underline{m}-1}$.

Proof. When the partnership sells the asset, all neighboring independent traders in rows $r \in \{\underline{m}, \ldots, \bar{m}+1\}$ bid ℓ. High-cost traders in row $\underline{m}-1$ bid ℓ and low-cost traders in that row bid $\tilde{b}_{\underline{m}-1}$. Therefore, the expected sale price is $\tilde{\delta}_{\underline{m}-1}\tilde{b}_{\underline{m}-1}$. Given Remark B.2, further trading profits are not possible once the asset reaches row $\underline{m}-1$. Therefore, $\tilde{\nu}_{\mathbf{m}} = \tilde{\delta}_{\underline{m}-1}\tilde{b}_{\underline{m}-1}$. □

Lemma B.2. *Take an arbitrary trading history where it is a commonly held belief among independent traders that the partnership has high trading costs with probability 1. Given* σ^*_{-i}, *the asset's expected resale value to independent trader* i *in row* $r \geq \underline{m}$ *is* $\tilde{\nu}_r = \tilde{\delta}_{r-1}\tilde{b}_{r-1} = \prod_{k=\underline{m}}^{r} \tilde{\delta}_{k-1}\tilde{b}_{\underline{m}-1}$.

Proof. The proof is by induction on r.

Base Case When i in row $r = \underline{m}$ sells the asset, all neighboring independent traders in rows $\underline{m}$ and $\underline{m}+1$ bid ℓ. Likewise, the partnership bids ℓ. Thus, the expected selling price is $\tilde{\delta}_{\underline{m}-1}\tilde{b}_{\underline{m}-1}$. By Remark B.2, further trading profits are not anticipated once the asset reaches row $\underline{m}-1$. Therefore, $\tilde{\nu}_{\underline{m}} = \tilde{\delta}_{\underline{m}-1}\tilde{b}_{\underline{m}-1}$.

Induction Hypothesis ($\star$) The asset's expected resale value to an independent trader in row r' is $\tilde{\nu}_{r'} = \prod_{k=\underline{m}}^{r'} \tilde{\delta}_{k-1}\tilde{b}_{\underline{m}-1}$.

Inductive Step The induction hypothesis is true for $r' = \underline{m}$. Suppose that it is true for $r' = r-1$. We will verify that it is true for $r' = r$. When agent i in row r sells the asset, neighboring independent traders in rows r and $r+1$ bid ℓ and the partnership bids ℓ. Given σ^*_{r-1}, low-cost independent traders in row $r-1$ are expected to bid $\tilde{b}_{r-1} = \tilde{\nu}_{r-1}$. Thus, the expected selling price is $\tilde{\delta}_{r-1}\tilde{b}_{r-1}$. By an argument parallel to that establishing Lemma A1 (but assuming the partnership bids ℓ when it has a chance to bid), we conclude that once the asset reaches row $r-1$, further trading profits are not expected by agent i. Therefore, the asset's expected resale value is $\tilde{\nu}_r = \tilde{\delta}_{r-1}\tilde{b}_{r-1}$. By ($\star$), $\tilde{b}_{r-1} = \tilde{\nu}_{r-1} = \prod_{k=\underline{m}}^{r-1} \tilde{\delta}_{k-1}\tilde{b}_{\underline{m}-1}$. Thus, $\tilde{\nu}_r = \prod_{k=\underline{m}}^{r} \tilde{\delta}_{k-1}\tilde{b}_{\underline{m}-1}$. □

Lemma B.3. *Take an arbitrary trading history where it is a commonly held belief among independent traders that the partnership has low trading costs with probability 1. Given* σ^*_{-i}, *the asset's expected resale value to independent trader i in row $r \geq \underline{m}$ is*

$$\tilde{\nu}_r = \begin{cases} \tilde{\delta}_{\underline{m}-1}\tilde{b}_{\underline{m}-1} & r = \underline{m} \\ \prod_{k=\underline{m}+1}^{r} \tilde{\mu}_{k-1} \cdot \tilde{\delta}_{\underline{m}-1}\tilde{b}_{\underline{m}-1} & \underline{m}+1 \leq r \leq \bar{m}+1 \\ \prod_{k=\bar{m}+2}^{r} \tilde{\delta}_{k-1} \cdot \prod_{k=\underline{m}}^{\bar{m}} \tilde{\mu}_k \cdot \tilde{\delta}_{\underline{m}-1}\tilde{b}_{\underline{m}-1} & r \geq \bar{m}+2 \end{cases} . \qquad \text{(B.1)}$$

Proof. The proof is by induction on r.

Base Case When i in row $r = \underline{m}$ sells the asset, all neighboring independent traders in rows $\underline{m}$ and $\underline{m}+1$ bid ℓ. Likewise, the partnership bids ℓ. Thus, the expected selling price is $\tilde{\delta}_{\underline{m}-1}\tilde{b}_{\underline{m}-1}$. By Remark B.2, further trading profits are not expected once the asset reaches row $\underline{m}-1$. Therefore, $\tilde{\nu}_{\underline{m}} = \tilde{\delta}_{\underline{m}-1}\tilde{b}_{\underline{m}-1}$.

Induction Hypothesis ($\star$) The asset's expected resale value to an independent trader in row r' is $\tilde{\nu}_{r'}$ as defined in (B.1).

Inductive Step The induction hypothesis is true for $r' = \underline{m}$. Suppose that it is true for $r' = r-1$. We will verify that it is true for $r' = r$. When agent i in row r sells the asset, neighboring independent traders in rows r and $r+1$ bid ℓ. Given σ^*_{r-1} and ($\star$), low-cost traders in row $r-1$ bid $\tilde{\nu}_{r-1}$. There are two sub-cases:

1. Suppose $\underline{m}+1 \leq r \leq \bar{m}+1$. Then $\tilde{b}_{\mathbf{m}} = \tilde{\nu}_{\mathbf{m}}$. Since $\tilde{\nu}_{\mathbf{m}} \geq \tilde{\nu}_{r-1}$, the expected sale price is $\tilde{\mu}_{r-1}\tilde{\nu}_{r-1} = \prod_{k=\underline{m}+1}^{r} \tilde{\mu}_{k-1} \cdot \tilde{\delta}_{\underline{m}-1}\tilde{b}_{\underline{m}-1}$. Next we confirm that after selling the asset (for price $\tilde{\nu}_{r-1}$) trader i cannot earn further trading profits. Two types of continuation histories are relevant.

   (a) Suppose the partnership acquires the asset. When it sells it, all independent traders, except perhaps i, in rows $r' \geq \underline{m}$ bid ℓ given σ^*_{-i}. Low-cost traders in row $\underline{m}-1$ bid $\tilde{b}_{\underline{m}-1}$. If a trader in row $\underline{m}-1$ acquires the asset, i will not have a further chance to purchase it. Suppose instead that i purchases the asset from the partnership by placing a competitive bid. The price paid by i can be one of two values. If the price is zero, then $\tilde{\mu}'_{\underline{m}-1}\tilde{b}_{\underline{m}-1} = 0 \implies \tilde{\delta}'_{\underline{m}-1}\tilde{b}_{\underline{m}-1} = 0$ where $\tilde{\mu}'_{\underline{m}-1}$ and $\tilde{\delta}'_{\underline{m}-1}$ correspond to updated beliefs (if applicable). Thus, when agent i sells the asset, all neighbors who submit a competitive bid will bid at most zero, precluding any further trading profits. If, however, the price is $\tilde{b}_{\underline{m}-1}$, then there is at least one low-cost trader in row $\underline{m}-1$. Despite this fact, when i sells the asset, all neighbors will bid at most $\tilde{b}_{\underline{m}-1}$ given σ^*_{-i}. Again, this implies profitable resale is not possible.

(b) Suppose an independent trader in row $r-1$ acquires the asset. If the asset reaches row $\underline{m}-1$ before being available to agent i again for purchase, he will not be able to profit further given the specified strategy. Similarly, if the asset reaches the partnership, then by the previous part, trader i will also not be able to profit further. Thus, suppose i purchases the asset directly from an independent trader in row $r-1$. If $r = \underline{m}+1$, then the reasoning of Lemma A1 applies thereby precluding further trading profits for agent i. Suppose instead that $r \geq \underline{m}+2$. Since $\tilde{\nu}_{\mathbf{m}} \geq \tilde{\nu}_{r-2}$, agent i must pay $\tilde{\nu}_{\mathbf{m}}$ for the asset. $\tilde{\nu}_{\mathbf{m}}$ is also an upper bound on all bids submitted when i sells the asset. Thus, trader i is unable to earn a positive profit by purchasing the asset from an independent trader in row $r-1$.

As additional trading profits are not possible, the agent's expected resale value equals the expected price from the original sale, $\tilde{\nu}_r = \prod_{k=\underline{m}+1}^{r} \tilde{\mu}_{k-1} \cdot \tilde{\delta}_{\underline{m}-1}\tilde{b}_{\underline{m}-1}$.

2. If $r \geq \bar{m}+2$, then the partnership is not directly relevant. Thus, the conclusion follows from Lemma A2 via induction and $\tilde{b}_{\bar{m}+1} = \tilde{\nu}_{\bar{m}+1}$.

□

Lemma B.4. *Take an arbitrary trading history in which agents in row $\bar{m}$ and the partnership have not yet placed any bids. Suppose a trader in row $\bar{m}+1$ sells the asset. Given σ^*,*

1. The expected bid of an independent, low-cost trader in row $\bar{m}$ is

$$\tilde{b}_{\bar{m}} = \prod_{k=\underline{m}}^{\bar{m}} \delta(n_{k-1} - m_{k-1})\tilde{b}_{\underline{m}-1}. \tag{B.2}$$

2. The expected bid of a low-cost partnership is $\tilde{b}_{\mathbf{m}} = \delta(n_{\underline{m}-1})\tilde{b}_{\underline{m}-1}$.

Thus, $\tilde{b}_{\bar{m}} < \tilde{b}_{\mathbf{m}}$.

Proof. Given the asset's trading history, agents hold their prior beliefs concerning the types of independent traders in row $r < \bar{m}$. Thus, $\tilde{\delta}_r = \delta(n_r - m_r)$. Applying Lemmas B.2 and B.3 gives the conclusion. □

B.5 Proof of Theorem B.1

We divide the proof of Theorem B.1 into cases corresponding to the defined strategy profile.

Case 1: Independent Trader i in Row $r \leq \underline{m} - 1$

The bidding problem faced by an independent trader in row $r \leq \underline{m} - 1$ is identical to that of a trader in a trading network without a partnership and Theorem 1 applies. From the proof of Theorem 1, we note that $\tilde{b}_{\underline{m}-1} \leq v$, which we employ below.

Case 2: The Partnership

Consider an arbitrary trading history and suppose that the asset is being sold by an independent trader in row $\underline{m} - 1$ or $\underline{m}$. Given $\sigma^*_{-\mathbf{m}}$ and Remark B.2, the maximal additional trading profit that the partnership can earn is zero. Therefore, the bid ℓ is optimal.

Suppose instead that the asset is sold by an independent trader in row $r \in \{\underline{m}+1, \ldots, \bar{m}+1\}$. To verify the optimality of $\sigma^*_{\mathbf{m}}$, we proceed by induction on r.

Base Case Suppose that the asset is currently held by independent trader i in row $\underline{m} + 1$. The asset's expected resale value to the partnership is $\tilde{\nu}_{\mathbf{m}} = \tilde{\delta}_{\underline{m}-1}\tilde{b}_{\underline{m}-1}$ (Lemma B.1). Since $\tilde{b}_{\underline{m}-1} \leq v$, it follows that $\tilde{\nu}_{\mathbf{m}} \leq v$. Thus, ℓ is an optimal bid for a high-cost partnership.

Next we confirm that $\tilde{\nu}_{\mathbf{m}}$ is an optimal bid for the partnership if it has low trading cost. Given the asset's trading history, an independent trader in row $\underline{m}$ believes that either the partnership has high trading costs or low trading costs. Lemmas B.2 and B.3 imply that this trader's expected resale value is $\tilde{\nu}_{\underline{m}} = \tilde{\delta}_{\underline{m}-1}\tilde{b}_{\underline{m}-1}$. Given $\sigma^*_{-\mathbf{m}}$, $\tilde{b}_{\underline{m}} = \tilde{\nu}_{\underline{m}} = \tilde{\nu}_{\mathbf{m}}$.

- If the partnership bids $\tilde{\nu}_{\mathbf{m}}$ (or more), the partnership realizes a trading profit only if all independent traders in row $\underline{m}$ bid ℓ. Thus, its expected trading profits are $(1 - \tilde{\mu}_{\underline{m}})\tilde{\nu}_{\mathbf{m}}$.

- If the partnership places a competitive bid strictly less than $\tilde{\nu}_{\mathbf{m}}$, its expected trading profits are also $(1 - \tilde{\mu}_{\underline{m}})\tilde{\nu}_{\mathbf{m}}$ (it wins only if all others bid ℓ given the prescribed behavior of independent traders).

- If the partnership bids ℓ, either trade breaks down or the asset is transferred to a low-cost, independent trader in row $\underline{m}$. By Remark B.2, the partnership cannot earn further trading profits given $\sigma^*_{\mathbf{m}}$. Thus, the partnership's trading profit is zero.

Hence, the partnership cannot improve upon its payoff from the bid $\tilde{\nu}_{\mathbf{m}}$.

Induction Hypothesis ($\star$) Whenever the asset is sold by an independent trader in row $r' \in \{\underline{m}+1, \ldots, k\}$, it is optimal for a low-cost partnership to bid $\tilde{\nu}_{\mathbf{m}} = \tilde{\delta}_{\underline{m}-1}\tilde{b}_{\underline{m}-1}$ and for a high-cost partnership to bid ℓ.

Inductive Step The base case ($k = \underline{m}+1$) satisfies the induction hypothesis. Assume ($\star$) is true for $k = r-1$. We will show that it is true for $k = r$.

Take an arbitrary trading history and suppose that the asset is being sold by independent trader i in row r. There are three cases depending on the trading history.

1. Suppose the partnership bid ℓ in the first auction in which it bid. Thus, independent traders believe that the partnership has a high trading cost with probability 1. Given $\sigma^*_{-\mathbf{m}}$ and Lemma B.2, $\tilde{b}_{r-1} = \prod_{k=\underline{m}}^{r-1} \tilde{\delta}_{k-1}\tilde{b}_{\underline{m}-1}$ and $\tilde{b}_r = \tilde{b}_{r+1} = \ell$. Note that $\tilde{\nu}_{\mathbf{m}} \geq \tilde{b}_{r-1}$.

   - If the partnership bids more than $\tilde{b}_{r-1}$, it may earn a profit under two circumstances. With probability $1 - \tilde{\nu}_{r-1}$ all independent traders bid ℓ and the partnership pays zero. With probability $\tilde{\nu}_{r-1}$, at least one independent trader bids $\tilde{b}_{r-1}$ which becomes the price paid by the partnership. Hence, its expected trading profit is

$$\begin{aligned}(1 - \tilde{\mu}_{r-1})\tilde{\nu}_{\mathbf{m}} + \tilde{\mu}_{r-1}(\tilde{\nu}_{\mathbf{m}} - \tilde{b}_{r-1}) &= \tilde{\nu}_{\mathbf{m}} - \tilde{\mu}_{r-1}\tilde{b}_{r-1} \\ &= \tilde{\nu}_{\mathbf{m}} - \tilde{\mu}_{r-1}\tilde{\delta}_{r-2}\prod_{k=\underline{m}}^{r-2} \tilde{\delta}_{k-1}\tilde{b}_{\underline{m}-1}.\end{aligned}$$

   - If the partnership places a competitive bid less that $\tilde{b}_{r-1}$, its expected trading profit is

$$(1 - \tilde{\mu}_{r-1}) \underbrace{\tilde{\nu}_{\mathbf{m}}}_{(A)} + \tilde{\mu}_{r-1} \left(\underbrace{(1 - \tilde{\mu}_{r-2})\tilde{\nu}_{\mathbf{m}} + \tilde{\mu}_{r-2}(\tilde{\nu}_{\mathbf{m}} - \tilde{b}_{r-2})}_{(B)} \right). \tag{B.3}$$

   (B.3) has two components. (A) With probability $(1-\tilde{\mu}_{r-1})$ all independent traders in row $r-1$ have high trading costs (and bid ℓ) and the partnership acquires the asset at zero cost. (B) With probability $\tilde{\mu}_{r-1}$ there is at least one low-cost independent trader in row $r-1$ who acquires the asset. In this case, the partnership has the opportunity to purchase the asset again when that agent sells it. By ($\star$), it is optimal for the partnership to bid $\tilde{\nu}_{\mathbf{m}}$ in that contingency. The bracketed term is the partnership's resulting expected profit. Collecting terms in (B.3) gives

$$\tilde{\nu}_{\mathbf{m}} - \tilde{\mu}_{r-1}\tilde{\mu}_{r-2}\tilde{b}_{r-2} = \tilde{\nu}_{\mathbf{m}} - \tilde{\mu}_{r-1}\tilde{\mu}_{r-2}\prod_{k=\underline{m}}^{r-2}\tilde{\delta}_{k-1}\tilde{b}_{\underline{m}-1}.$$

- If the partnership bids ℓ, its expected trading profit is

$$\tilde{\mu}_{r-1}\tilde{\nu}_{\mathbf{m}} - \tilde{\mu}_{r-1}\tilde{\mu}_{r-2}\tilde{b}_{r-2} = \tilde{\mu}_{r-1}\tilde{\nu}_{\mathbf{m}} - \tilde{\mu}_{r-1}\tilde{\mu}_{r-2}\prod_{k=\underline{m}}^{r-2}\tilde{\delta}_{k-1}\tilde{b}_{\underline{m}-1}. \tag{B.4}$$

  The derivation of (B.4) mirrors the reasoning of the preceding case.

Since $\tilde{\delta}_{r-2} \leq \tilde{\mu}_{r-2}$, $\tilde{\nu}_{\mathbf{m}}$ is an optimal bid for a low-cost partnership. If it has high cost, ℓ is an optimal bid as all competitive bids yield an expected profit less than v.

2. Suppose the partnership placed a competitive bid in the first auction in which it bid. Thus, independent traders believe it has low trading costs with probability 1. Given $\sigma^*_{-\mathbf{m}}$ and Lemma B.3,

$$\tilde{b}_{r-1} = \prod_{k=\underline{m}+1}^{r-1}\tilde{\mu}_{k-1}\tilde{\delta}_{\underline{m}-1}\tilde{b}_{\underline{m}-1}$$

and $\tilde{b}_r = \tilde{b}_{r+1} = \ell$. Note that $\tilde{\nu}_{\mathbf{m}} \geq \tilde{b}_{r-1}$.

- If the partnership bids more than $\tilde{b}_{r-1}$, its expected trading profit is

$$\begin{aligned}&(1-\tilde{\mu}_{r-1})\tilde{\nu}_{\mathbf{m}} + \tilde{\mu}_{r-1}(\tilde{\nu}_{\mathbf{m}} - \tilde{b}_{r-1})\\ &= \tilde{\nu}_{\mathbf{m}} - \tilde{\mu}_{r-1}\tilde{b}_{r-1}\\ &= \tilde{\nu}_{\mathbf{m}} - \tilde{\mu}_{r-1}\prod_{k=\underline{m}+1}^{r-1}\tilde{\mu}_{k-1}\tilde{\delta}_{\underline{m}-1}\tilde{b}_{\underline{m}-1}.\end{aligned}$$

- If the partnership places a competitive bid less than $\tilde{b}_{r-1}$, its expected trading profit is

$$\begin{aligned}&(1-\tilde{\mu}_{r-1})\tilde{\nu}_{\mathbf{m}}+\tilde{\mu}_{r-1}\Big((1-\tilde{\mu}_{r-2})\tilde{\nu}_{\mathbf{m}}+\tilde{\mu}_{r-2}(\tilde{\nu}_{\mathbf{m}}-\tilde{b}_{r-2})\Big)\\&=\tilde{\nu}_{\mathbf{m}}-\tilde{\mu}_{r-1}\tilde{\mu}_{r-2}\tilde{b}_{r-2}\\&=\tilde{\nu}_{\mathbf{m}}-\tilde{\mu}_{r-1}\tilde{\mu}_{r-2}\prod_{k=\underline{m}+1}^{r-2}\tilde{\mu}_{k-1}\tilde{\delta}_{\underline{m}-1}\tilde{b}_{\underline{m}-1}\end{aligned}$$

The derivation of the preceding expressions mirrors that of the analogous situation in case 1 above.

- If the partnership bids ℓ, its expected trading profit is

$$\tilde{\mu}_{r-1}\tilde{\nu}_{\mathbf{m}}-\tilde{\mu}_{r-1}\tilde{\mu}_{r-2}\prod_{k=\underline{m}+1}^{r-2}\tilde{\mu}_{k-1}\tilde{\delta}_{\underline{m}-1}\tilde{b}_{\underline{m}-1}.$$

Comparing the above expressions, by inspection we can conclude that $\tilde{\nu}_{\mathbf{m}}$ is an optimal bid for the partnership if it has low trading cost. The bid ℓ is optimal if it has high trading costs.

3. Suppose the partnership has not placed any bids. Therefore, $r=\bar{m}+1$ and all agents in rows $k\leq\bar{m}$ have not yet bid. Thus, $\tilde{\mu}_k=\mu(n_k-m_k)$ and $\tilde{\delta}_k=\delta(n_k-m_k)$ for all $k\leq\bar{m}$. From Lemma B.4, $\tilde{b}_{\bar{m}}=\prod_{k=\underline{m}}^{\bar{m}-1}\delta(n_k-m_k)\cdot\delta(n_{\underline{m}-1})\tilde{b}_{\underline{m}-1}$. Clearly, $\tilde{\nu}_{\mathbf{m}}>\tilde{b}_{\bar{m}}$.

   - If the partnership bids $\tilde{b}_{\bar{m}}$ (or more), its expected payoff is

$$\tilde{\nu}_{\mathbf{m}}-\tilde{\mu}_{\bar{m}}\tilde{b}_{\bar{m}}=\tilde{\nu}_{\mathbf{m}}-\tilde{\mu}_{\bar{m}}\prod_{k=\underline{m}}^{\bar{m}-1}\tilde{\delta}_k\cdot\tilde{\delta}_{\underline{m}-1}\tilde{b}_{\underline{m}-1}. \tag{B.5}$$

   - If the partnership places a competitive bid less than $\tilde{b}_{\bar{m}}$, then its expected trading profit is

$$(1-\tilde{\mu}_{\bar{m}})\tilde{\nu}_{\mathbf{m}}+\tilde{\mu}_{\bar{m}}\Big((1-\tilde{\mu}_{\bar{m}-1})\tilde{\nu}_{\mathbf{m}}+\tilde{\mu}_{\bar{m}-1}(\tilde{\nu}_{\mathbf{m}}-\tilde{b}'_{\bar{m}-1})\Big)$$

   where, given $\sigma^*_{-\mathbf{m}}$ and the assumed evolution of agents' beliefs,

$$\tilde{b}'_{\bar{m}-1}=\prod_{k=\underline{m}+1}^{\bar{m}-1}\tilde{\mu}_{k-1}\cdot\tilde{\delta}_{\underline{m}-1}\tilde{b}_{\underline{m}-1}.$$

Substituting $\tilde{b}'_{\bar{m}-1}$ into the above expression gives an expected profit of

$$\tilde{\nu}_{\mathbf{m}} - \tilde{\mu}_{\bar{m}}\tilde{\mu}_{\bar{m}-1} \prod_{k=\underline{m}+1}^{\bar{m}-1} \tilde{\mu}_{k-1} \cdot \tilde{\delta}_{\underline{m}-1}\tilde{b}_{\underline{m}-1}. \tag{B.6}$$

- If the partnership bids ℓ, then its expected trading profit is

$$\tilde{\mu}_{\bar{m}} \left((1 - \tilde{\mu}_{\bar{m}-1})\tilde{\nu}_{\mathbf{m}} + \tilde{\mu}_{\bar{m}-1}(\tilde{\nu}_{\mathbf{m}} - \tilde{b}''_{\bar{m}-1}) \right)$$

  where, given $\sigma^*_{-\mathbf{m}}$ and the assumed evolution of traders' beliefs, $\tilde{b}''_{\bar{m}-1} = \prod_{k=\underline{m}+1}^{\bar{m}-1} \tilde{\delta}_{k-1} \cdot \tilde{\delta}_{\underline{m}-1}\tilde{b}_{\underline{m}-1}$. Substituting $\tilde{b}''_{\bar{m}-1}$ into the above expression gives an expected profit of

$$\tilde{\mu}_{\bar{m}}\tilde{\nu}_{\mathbf{m}} - \tilde{\mu}_{\bar{m}}\tilde{\mu}_{\bar{m}-1} \prod_{k=\underline{m}+1}^{\bar{m}-1} \tilde{\delta}_{k-1} \cdot \tilde{\delta}_{\underline{m}-1}\tilde{b}_{\underline{m}-1}. \tag{B.7}$$

Since $\tilde{\delta}_k \leq \tilde{\mu}_k$, comparing (B.5) to (B.6) and (B.7) shows that $\tilde{\nu}_{\mathbf{m}}$ is an optimal bid if the partnership has low trading cost. Else, since the above expressions are all less than v, ℓ is the optimal bid if the partnership has high trading costs.

The three cases considered above exhaust all possibilities, thereby verifying the claim for $k = r$.

Case 3: Independent Trader i in Row $\underline{m} \leq r \leq \bar{m} - 1$

There are two cases depending on the asset's trading history. Either the partnership bid ℓ in the first auction in which it bid or it placed a competitive bid. The asset's expected resale values in these cases are $\tilde{\nu}_r = \prod_{k=\underline{m}}^{r} \tilde{\delta}_{k-1}\tilde{b}_{\underline{m}-1}$ and $\tilde{\nu}_r = \prod_{k=\underline{m}+1}^{r} \tilde{\mu}_{k-1} \cdot \tilde{\delta}_{\underline{m}-1}\tilde{b}_{\underline{m}-1}$, respectively. The difference stems from the different anticipated behavior of the partnership in future auctions.

1. Suppose the asset is sold by another independent trader in row $r-1$ or r. By the same reasoning used to establish Lemma A1, i cannot earn further trading profits conditional on the asset's location. Thus, ℓ is an optimal bid.

2. Suppose the asset is sold by the partnership. If at least one row$-(\underline{m} - 1)$ trader bids $\tilde{b}_{\underline{m}-1}$, then agent i must pay at least $\tilde{b}_{\underline{m}-1}$. However, the expected resale value is bounded above by $\tilde{b}_{\underline{m}-1}$; therefore, a profit cannot be earned. On the other hand,

if all row$-(\underline{m}-1)$ traders bids ℓ, the asset's expected resale value is zero thereafter; therefore, a profit cannot be earned. Consequently, ℓ is an optimal bid for i.

3. Suppose the asset is sold by an independent trader in row $r+1$. Given the specification of beliefs and σ^*_{-i}, an argument that is parallel to that confirming that "bidding one's valuation" is optimal in a second price auction (Vickrey, 1961) shows that $\tilde{\nu}_r$ is an optimal bid.

Case 4: Independent Trader in Row $r = \bar{m}$

If the partnership has already placed a bid in the asset's trading history, the analysis of case 3, above, applies.

Suppose the partnership has not placed any bids and the asset is held by a trader in row $r = \bar{m}+1$. From Lemma B.4, the bid of a low-cost independent trader in row $\bar{m}$ is

$$\tilde{b}_{\bar{m}} = \prod_{k=\underline{m}+1}^{\bar{m}} \delta(n_{k-1} - m_{k-1}) \cdot \delta(n_{\underline{m}-1})\tilde{b}_{\underline{m}-1}.$$

With this bid, i can win the auction only if the partnership bids ℓ. In this contingency, the expected resale value of the asset is $\tilde{\nu}_{\bar{m}} = \prod_{k=\underline{m}}^{\bar{m}} \tilde{\delta}_{k-1}\tilde{b}_{\underline{m}-1}$ where $\tilde{\delta}_k = \delta(n_k - m_k)$. Thus, trader i's expected profit is nonnegative.

To verify that $\tilde{b}_{\bar{m}}$ is an optimal bid we consider the three possible alternatives:

1. If trader i bids strictly less than $\tilde{\nu}_{\mathbf{m}}$, his expected trading profit is the same as from the bid $\tilde{b}_{\bar{m}}$ given σ^*_{-i}.

2. If trader i bids ℓ, his expected profit is zero. Conditional on the asset being purchased by the partnership or by another trader in row $\bar{m}$, further trading profits are not possible for trader i by the reasoning in the proofs of Lemmas B.2 and B.3.

3. If trader i bids $\tilde{\nu}_{\mathbf{m}}$ or more, one of three events may occur. If the partnership bids ℓ, then i's expected payoff is the same as if he had bid $\tilde{b}_{\bar{m}}$. If the partnership bids $\tilde{\nu}_{\mathbf{m}}$ and i receives the asset, then $\tilde{\nu}_{\mathbf{m}}$ is i's payment. His expected resale value is $\prod_{k=\underline{m}}^{\bar{m}-1} \tilde{\mu}_k \tilde{\delta}_{\underline{m}-1}\tilde{b}_{\underline{m}-1} < \tilde{\nu}_{\mathbf{m}}$. Thus, he earns negative profits in this contingency. Finally, if i fails to acquire the asset, his immediate payoff is zero and further trading profits are not possible. Thus, a bid of $\tilde{\nu}_{\mathbf{m}}$ or more is not more profitable for trader i than $\tilde{b}_{\bar{m}}$.

Noting the above cases, we see that it is optimal for a low-cost trader to bid $\tilde{b}_{\bar{m}}$. If the trader has high cost, ℓ is optimal as the expected resale value is bounded above by v.

Case 5: Independent Trader in Row $r \geq \bar{m} + 1$

The bidding problem faced by an independent trader in row $r \geq m+1$ is identical to that of a trader in a trading network without a partnership and the argument of Theorem 1 applies.

The following appendix is intended for online publication only.

C Variation Measures and the Bullwhip Effect

This section complements our discussion of the bullwhip effect, which is a stylized fact observed in many supply chain networks concerning the variability of demand (Lee et al., 1997a).

Demand Variation When an agent in row $r+1$ sells the asset, either there is demand for the asset or there is no demand, i.e. it is binary. The expected demand is thus $\mu(n_r)$, which equals the probability that at least one agent in row r places a competitive bid in equilibrium. Thus, the standard deviation of demand is $\sqrt{(1-\mu(n_r))\mu(n_r)}$. Dividing the standard deviation by the expected value gives the term of interest:

$$CVD_r(\mathbf{n}) = \frac{\text{Standard Deviation of Demand}}{\text{Expected Demand}} = \frac{\sqrt{(1-\mu(n_r))\mu(n_r)}}{\mu(n_r)} = \sqrt{\frac{1}{\mu(n_r)} - 1}.$$

Price Variation When an agent in row $r+1$ sells the asset to an agent in row r, the expected sales price is $\nu_{r+1} = \prod_{k=1}^{r} \delta(n_k)$. A simple calculation shows that the standard deviation of that sales price equals

$$\sqrt{\prod_{k=1}^{r} \delta(n_k) \cdot \prod_{k=1}^{r-1} \delta(n_k) \cdot (1 - \delta(n_r))}.$$

Dividing the standard deviation by the expected value gives the term of interest:

$$CVP_r(\mathbf{n}) = \frac{\text{Standard Deviation of Price}}{\text{Expected Price}} = \sqrt{\frac{1}{\delta(n_r)} - 1}.$$

References

Acemoglu, D., Carvalho, V. M., Ozdaglar, A., and Tahbaz-Salehi, A. (2012). The network origins of aggregate fluctuations. *Econometrica*, 80(5):1977–2016.

Arrow, K. J. (1975). Vertical integration and communication. *The Bell Journal of Economics*, 6(1):173–183.

Bala, V. and Goyal, S. (2000). A noncooperative model of network formation. *Econometrica*, 68(5):1181–1229.

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. (2012a). Bilateral trading in networks. Mimeo.

Condorelli, D. and Galeotti, A. (2012b). Endogenous trading networks. Mimeo.

Corominas-Bosch, M. (2004). Bargaining in a network of buyers and sellers. *Journal of Economic Theory*, 115(1):35–77.

Economides, N. (1996). The economics of networks. *International Journal of Industrial Organization*, 14(6):673–699.

Elliott, M. (2013). Inefficiencies in networked markets. Mimeo.

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

Gale, D. and Shapley, L. S. (1962). College admissions and the stability of marriage. *The American Mathematical Monthly*, 69(1):9–15.

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

Gale, D. M., Kariv, S., Kotowski, M. H., and Leister, C. M. (2012). Work in Progress.

Gary-Bobo, R. J. (1990). On the existence of equilibrium points in a class of asymmetric market entry games. *Games and Economic Behavior*, 2(3):239–246.

Gofman, M. (2011). A network-based analysis of over-the-counter markets. Mimeo.

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

Jackson, M. O. and Wolinsky, A. (1996). A strategic model of social and economic networks. *Journal of Economic Theory*, 71(1):44–74.

Kranton, R. E. and Minehart, D. F. (2000). Networks versus vertical integration. *RAND Journal of Economics*, 31(3):570–601.

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

Kremer, M. (1993). The O-ring theory of economic development. *Quarterly Journal of Economics*, 108(3):551–575.

Lee, H. L., Padmanabhan, V., and Whang, S. (1997a). The bullwhip effect in supply chains. *MIT Sloan Management Review*, 38(3).

Lee, H. L., Padmanabhan, V., and Whang, S. (1997b). Information distortion in a supply chain: The bullwhip effect. *Management Science*, 43(4):546–558.

Manea, M. (2011). Bargaining in stationary networks. *American Economic Review*, 101(5):2042–80.

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

Mankiw, N. G. and Whinston, M. D. (1986). Free entry and social inefficiency. *RAND Journal of Economics*, 17(1):48–58.

Myerson, R. B. (1981). Optimal auction design. *Mathematics of Operations Research*, 6(1):58–73.

Nagurney, A. and Qiang, Q. (2009). *Fragile Networks: Identifying Vulnerabilities and Synergies in an Uncertain World*. John Wiley & Sons, Hoboken, NJ.

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

Patil, H. (2011). Buyer-seller networks with demand shocks and intermediation. *Review of Economic Design*, 15(2):121–145.

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

Siedlarek, J.-P. (2012). Intermediation in networks. Mimeo.

Spulber, D. F. (1996). Market microstructure and intermediation. *Journal of Economic Perspectives*, 10(3):135–152.

Vickrey, W. (1961). Counterspeculation, auctions, and competitive sealed tenders. *Journal of Finance*, 16(1):8–37.

Wright, R. and Wong, Y.-Y. (2014). Buyers, sellers, and middlement: Variations on search-theoretic themes. *International Economic Review*, 55(2):375–397.

Zheng, C. Z. (2002). Optimal auction with resale. *Econometrica*, 70(6):2197–2224.